$270.00

Lab Reference for

Building Scalable Networks with

Advanced Routing Configurations

www.hiteksolutions.com

- write erase
- reload

Kamoa
AskKamoa@hotmail.com

Basic OSPF router configuration in a single area

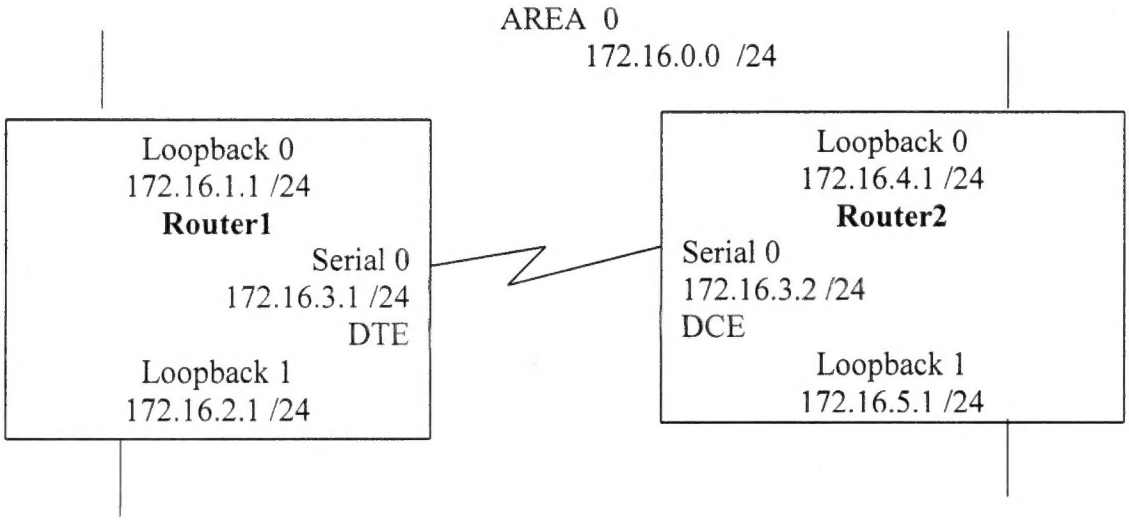

Config t
Hostname router1
Ip routing
Int loopback0
Ip address 172.16.1.1 255.255.255.0
Int loopback1
Ip address 172.16.2.1 255.255.255.0
Int serial0
Ip address 172.16.3.1 255.255.255.0
No shutdown

Router ospf 10
Network 172.16.1.1 0.0.0.0 area 0
Network 172.16.2.1 0.0.0.0 area 0
Network 172.16.3.1 0.0.0.0 area 0

OR ANOTHER WAY

Router ospf 10
Network 172.16.0.0 0.0.255.255 area 0

Config t
hostname router2
ip routing
int loopback0
ip address 172.16.4.1 255.255.255.0
int loopback1
ip address 172.16.5.1 255.255.255.0
int serial0
ip address 172.16.3.2 255.255.255.0
no shutdown
Clock rate 56000

router ospf 10
network 172.16.4.1 0.0.0.0 area 0
network 172.16.5.1 0.0.0.0 area 0
network 172.16.3.2 0.0.0.0 area 0

OR ANOTHER WAY

router ospf 10
network 172.16.0.0 0.0.255.255 area 0

172.16.0.0. 0.0.0.0 0.0.0.0. area 0
include "all" interfaces

```
router1#config t
Enter configuration commands, one per line.  End with CNTL/Z.
router1(config)#router ?
  bgp        Border Gateway Protocol (BGP)
  egp        Exterior Gateway Protocol (EGP)
  eigrp      Enhanced Interior Gateway Routing Protocol (EIGRP)
  igrp       Interior Gateway Routing Protocol (IGRP)
  isis       ISO IS-IS
  iso-igrp   IGRP for OSI networks
  mobile     Mobile routes
  odr        On Demand stub Routes
  ospf       Open Shortest Path First (OSPF)
  rip        Routing Information Protocol (RIP)
  static     Static routes

router1(config)#router ospf ?
  <1-65535>  Process ID

router1(config)#router ospf 10
router1(config-router)#
router1(config-router)#?
Router configuration commands:
  area                  OSPF area parameters
  default               Set a command to its defaults
  default-information   Control distribution of default information
  default-metric        Set metric of redistributed routes
  distance              Define an administrative distance
  distribute-list       Filter networks in routing updates
  exit                  Exit from routing protocol configuration mode
  help                  Description of the interactive help system
  maximum-paths         Forward packets over multiple paths
  neighbor              Specify a neighbor router
  network               Enable routing on an IP network
  no                    Negate a command or set its defaults
  ospf                  OSPF specific commands
  passive-interface     Suppress routing updates on an interface
  redistribute          Redistribute information from another routing protocol
  summary-address       Configure IP address summaries
  timers                Adjust routing timers

router1(config-router)#network 172.16.1.1 0.0.0.0 area 0
router1(config-router)#network 172.16.2.1 0.0.0.0 area 0
router1(config-router)#network 172.16.3.1 0.0.0.0 area 0
router1(config-router)#^Z

router1#
router1#sh run
Building configuration...

Current configuration:
!
version 11.2
no service udp-small-servers
no service tcp-small-servers
!
hostname router1
!
enable secret 5 $1$wgkP$I68COYt3xL9sTiuJNOmCJ0
enable password 12345
```

```
!
!
interface Loopback0
 ip address 172.16.1.1 255.255.255.0
!
interface Loopback1
 ip address 172.16.2.1 255.255.255.0
!
interface Ethernet0
 no ip address
 shutdown
!
interface Serial0
 ip address 172.16.3.1 255.255.255.0
!
router ospf 10
 network 172.16.1.1 0.0.0.0 area 0
 network 172.16.2.1 0.0.0.0 area 0
 network 172.16.3.1 0.0.0.0 area 0
!
no ip classless
!
line con 0
line vty 0 4
 password 12345
 login
!
end

router1#sh ip route
Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, * - candidate default
       U - per-user static route, o - ODR

Gateway of last resort is not set

     172.16.0.0/16 is variably subnetted, 5 subnets, 2 masks
O       172.16.5.1/32 [110/65] via 172.16.3.2, 00:03:56, Serial0
O       172.16.4.1/32 [110/65] via 172.16.3.2, 00:03:56, Serial0
C       172.16.1.0/24 is directly connected, Loopback0
C       172.16.2.0/24 is directly connected, Loopback1
C       172.16.3.0/24 is directly connected, Serial0
router1#

router1#sh ip route ospf
     172.16.0.0/16 is variably subnetted, 5 subnets, 2 masks
O       172.16.5.1/32 [110/65] via 172.16.3.2, 00:04:19, Serial0
O       172.16.4.1/32 [110/65] via 172.16.3.2, 00:04:19, Serial0
router1#
router1#

router1#sh ip protocol
Routing Protocol is "ospf 10"
  Sending updates every 0 seconds
  Invalid after 0 seconds, hold down 0, flushed after 0
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set
```

```
  Redistributing: ospf 10
  Routing for Networks:
    172.16.1.1/32
    172.16.2.1/32
    172.16.3.1/32
  Routing Information Sources:
    Gateway         Distance      Last Update
    172.16.5.1           110      00:16:29
  Distance: (default is 110)

router1#

router1#sh ip ospf ?
  <1-4294967295>       Process ID number
  border-routers       Border and Boundary Router Information
  database             Database summary
  interface            Interface information
  neighbor             Neighbor list
  request-list         Link state request list
  retransmission-list  Link state retransmission list
  summary-address      Summary-address redistribution Information
  virtual-links        Virtual link information
  <cr>

router1#sh ip ospf
 Routing Process "ospf 10" with ID 172.16.2.1
 Supports only single TOS(TOS0) routes
 SPF schedule delay 5 secs, Hold time between two SPFs 10 secs
 Number of DCbitless external LSA 0
 Number of DoNotAge external LSA 0
 Number of areas in this router is 1. 1 normal 0 stub 0 nssa
    Area BACKBONE(0)
      Number of interfaces in this area is 3
      Area has no authentication
      SPF algorithm executed 2 times
      Area ranges are
      Link State Update Interval is 00:30:00 and due in 00:12:11
      Link State Age Interval is 00:20:00 and due in 00:02:10
      Number of DCbitless LSA 1
      Number of indication LSA 0
      Number of DoNotAge LSA 0

router1#sh ip ospf 10
 Routing Process "ospf 10" with ID 172.16.2.1
 Supports only single TOS(TOS0) routes
 SPF schedule delay 5 secs, Hold time between two SPFs 10 secs
 Number of DCbitless external LSA 0
 Number of DoNotAge external LSA 0
 Number of areas in this router is 1. 1 normal 0 stub 0 nssa
    Area BACKBONE(0)
      Number of interfaces in this area is 3
      Area has no authentication
      SPF algorithm executed 2 times
      Area ranges are
      Link State Update Interval is 00:30:00 and due in 00:11:58
      Link State Age Interval is 00:20:00 and due in 00:01:58
      Number of DCbitless LSA 1
      Number of indication LSA 0
      Number of DoNotAge LSA 0
```

```
router1#sh ip ospf database

       OSPF Router with ID (172.16.2.1) (Process ID 10)


            Router Link States (Area 0)

Link ID         ADV Router      Age         Seq#       Checksum Link count
172.16.2.1      172.16.2.1      1091        0x80000002 0x350B   4
172.16.5.1      172.16.5.1      1093        0x80000003 0x1A3B   4
router1#
router1#sh ip ospf database router

       OSPF Router with ID (172.16.2.1) (Process ID 10)


            Router Link States (Area 0)

  LS age: 1108
  Options: (No TOS-capability, DC)
  LS Type: Router Links
  Link State ID: 172.16.2.1
  Advertising Router: 172.16.2.1
  LS Seq Number: 80000002
  Checksum: 0x350B
  Length: 72
   Number of Links: 4

    Link connected to: another Router (point-to-point)
     (Link ID) Neighboring Router ID: 172.16.5.1
     (Link Data) Router Interface address: 172.16.3.1
      Number of TOS metrics: 0
       TOS 0 Metrics: 64

    Link connected to: a Stub Network
     (Link ID) Network/subnet number: 172.16.3.0
     (Link Data) Network Mask: 255.255.255.0
      Number of TOS metrics: 0
       TOS 0 Metrics: 64

    Link connected to: a Stub Network
     (Link ID) Network/subnet number: 172.16.2.1
     (Link Data) Network Mask: 255.255.255.255
      Number of TOS metrics: 0
       TOS 0 Metrics: 1

    Link connected to: a Stub Network
     (Link ID) Network/subnet number: 172.16.1.1
     (Link Data) Network Mask: 255.255.255.255
      Number of TOS metrics: 0
       TOS 0 Metrics: 1


  LS age: 1111
  Options: (No TOS-capability, No DC)
  LS Type: Router Links
  Link State ID: 172.16.5.1
  Advertising Router: 172.16.5.1
```

```
   LS Seq Number: 80000003
   Checksum: 0x1A3B
   Length: 72
    Number of Links: 4

     Link connected to: a Stub Network
      (Link ID) Network/subnet number: 172.16.5.1
      (Link Data) Network Mask: 255.255.255.255
       Number of TOS metrics: 0
        TOS 0 Metrics: 1

     Link connected to: a Stub Network
      (Link ID) Network/subnet number: 172.16.4.1
      (Link Data) Network Mask: 255.255.255.255
       Number of TOS metrics: 0
        TOS 0 Metrics: 1

     Link connected to: another Router (point-to-point)
      (Link ID) Neighboring Router ID: 172.16.2.1
      (Link Data) Router Interface address: 172.16.3.2
       Number of TOS metrics: 0
        TOS 0 Metrics: 64

     Link connected to: a Stub Network
      (Link ID) Network/subnet number: 172.16.3.0
      (Link Data) Network Mask: 255.255.255.0
       Number of TOS metrics: 0
        TOS 0 Metrics: 64

router1#

router1#sh ip ospf interface
Ethernet0 is administratively down, line protocol is down
   OSPF not enabled on this interface
Loopback0 is up, line protocol is up
  Internet Address 172.16.1.1/24, Area 0
  Process ID 10, Router ID 172.16.2.1, Network Type LOOPBACK, Cost: 1
  Loopback interface is treated as a stub Host
Loopback1 is up, line protocol is up
  Internet Address 172.16.2.1/24, Area 0
  Process ID 10, Router ID 172.16.2.1, Network Type LOOPBACK, Cost: 1
  Loopback interface is treated as a stub Host
Serial0 is up, line protocol is up
  Internet Address 172.16.3.1/24, Area 0
  Process ID 10, Router ID 172.16.2.1, Network Type POINT_TO_POINT, Cost: 64
  Transmit Delay is 1 sec, State POINT_TO_POINT,
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello due in 00:00:00
  Neighbor Count is 1, Adjacent neighbor count is 1
    Adjacent with neighbor 172.16.5.1
  Suppress hello for 0 neighbor(s)
router1#
router1#sh ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
172.16.5.1        1   FULL/  -        00:00:39     172.16.3.2      Serial0
router1#
router1#sh ip ospf neighbor detail
 Neighbor 172.16.5.1, interface address 172.16.3.2
    In the area 0 via interface Serial0
```

```
    Neighbor priority is 1, State is FULL
    Options 2
    Dead timer due in 00:00:36
router1#
router1#

router1#debug ip ospf ?
  adj             OSPF adjacency events
  events          OSPF events
  flood           OSPF flooding
  lsa-generation  OSPF lsa generation
  packet          OSPF packets
  retransmission  OSPF retransmission events
  spf             OSPF spf
  tree            OSPF database tree

router1#debug ip ospf events
OSPF events debugging is on
router1#
%LINEPROTO-5-UPDOWN: Line protocol on Interface Serial0, changed state to down
%LINK-3-UPDOWN: Interface Serial0, changed state to down
OSPF: Interface Serial0 going Down
OSPF: neighbor 172.16.2.1 is dead, state DOWN
OSPF: neighbor 172.16.5.1 is dead, state DOWN
OSPF: Tried to build Router LSA within MinLSInterval
%LINK-3-UPDOWN: Interface Serial0, changed state to up
OSPF: Interface Serial0 going Up
%LINEPROTO-5-UPDOWN: Line protocol on Interface Serial0, changed state to up
%LINEPROTO-5-UPDOWN: Line protocol on Interface Serial0, changed state to down
%LINK-3-UPDOWN: Interface Serial0, changed state to down
OSPF: Interface Serial0 going Down
OSPF: neighbor 172.16.2.1 is dead, state DOWN
%LINK-3-UPDOWN: Interface Serial0, changed state to up
OSPF: Interface Serial0 going Up
OSPF: Tried to build Router LSA within MinLSInterval
%LINEPROTO-5-UPDOWN: Line protocol on Interface Serial0, changed state to up
OSPF: Receive dbd from 172.16.5.1 seq 0x2203
OSPF: 2 Way Communication to neighbor 172.16.5.1
OSPF: send DBD packet to 172.16.3.2 seq 0x7AF
OSPF: NBR Negotiation Done  We are the SLAVE
OSPF: send DBD packet to 172.16.3.2 seq 0x2203
OSPF: Receive dbd from 172.16.5.1 seq 0x2204
OSPF: send DBD packet to 172.16.3.2 seq 0x2204
OSPF: Receive dbd from 172.16.5.1 seq 0x2205
OSPF: Exchange Done with neighbor 172.16.5.1
OSPF: Synchronized with neighbor 172.16.5.1, state:FULL
OSPF: Tried to build Router LSA within MinLSInterval
OSPF: send DBD packet to 172.16.3.2 seq 0x2205

router1#debug ip ospf adj
OSPF adjacency events debugging is on
router1#
router1#
%LINEPROTO-5-UPDOWN: Line protocol on Interface Serial0, changed state to down
%LINK-3-UPDOWN: Interface Serial0, changed state to down
OSPF: Interface Serial0 going Down
OSPF: neighbor 172.16.2.1 is dead, state DOWN
OSPF: Build router LSA, router ID 172.16.2.1
OSPF: neighbor 172.16.5.1 is dead, state DOWN
OSPF: Build router LSA, router ID 172.16.2.1
```

```
OSPF: Build router LSA, router ID 172.16.2.1
OSPF: Build router LSA, router ID 172.16.2.1
%LINK-3-UPDOWN: Interface Serial0, changed state to up
OSPF: Interface Serial0 going Up
OSPF: Build router LSA, router ID 172.16.2.1
%LINEPROTO-5-UPDOWN: Line protocol on Interface Serial0, changed state to up
%LINEPROTO-5-UPDOWN: Line protocol on Interface Serial0, changed state to down
%LINK-3-UPDOWN: Interface Serial0, changed state to down
OSPF: Interface Serial0 going Down
OSPF: neighbor 172.16.2.1 is dead, state DOWN
OSPF: Build router LSA, router ID 172.16.2.1
OSPF: Build router LSA, router ID 172.16.2.1
%LINK-3-UPDOWN: Interface Serial0, changed state to up
OSPF: Interface Serial0 going Up
OSPF: Build router LSA, router ID 172.16.2.1
%LINEPROTO-5-UPDOWN: Line protocol on Interface Serial0, changed state to up
OSPF: Build router LSA, router ID 172.16.2.1
OSPF: Receive dbd from 172.16.5.1 seq 0x1B61
OSPF: 2 Way Communication to neighbor 172.16.5.1
OSPF: send DBD packet to 172.16.3.2 seq 0x24E8
OSPF: NBR Negotiation Done  We are the SLAVE
OSPF: send DBD packet to 172.16.3.2 seq 0x1B61
OSPF: Receive dbd from 172.16.5.1 seq 0x1B62
OSPF: send DBD packet to 172.16.3.2 seq 0x1B62
OSPF: Receive dbd from 172.16.5.1 seq 0x1B63
OSPF: Exchange Done with neighbor 172.16.5.1
OSPF: Synchronized with neighbor 172.16.5.1, state:FULL
OSPF: Build router LSA, router ID 172.16.2.1
OSPF: send DBD packet to 172.16.3.2 seq 0x1B63
router1#undebug all
All possible debugging has been turned off
router1#
```

Basic OSPF router configuration for multiple area's

no backbone area → virtual Link

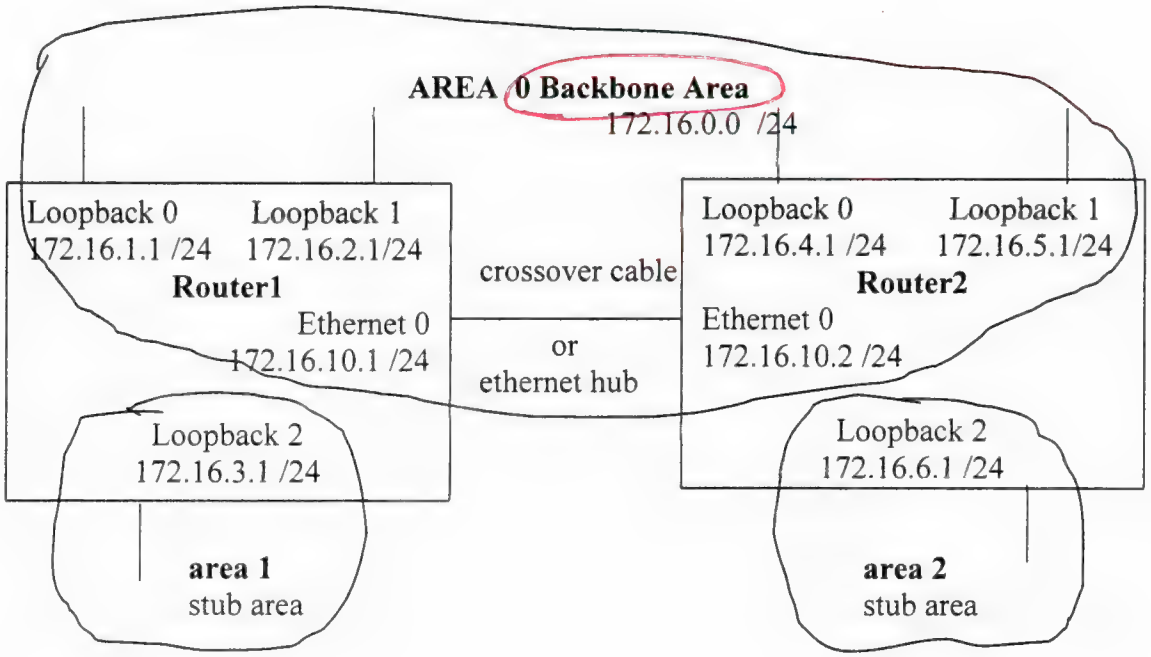

Config t
Hostname router1
Ip routing
Int loopback0
Ip address 172.16.1.1 255.255.255.0
Int loopback1
Ip address 172.16.2.1 255.255.255.0
Int loopback 2
Ip address 172.16.3.1 255.255.255.0
Int ethernert 0
172.16.10.1 255.255.255.0
no shutdown

Router ospf 20
Network 172.16.3.1 0.0.0.0 area 1
Network 172.16.0.0 0.0.255.255 area 0
Area 1 stub

Config t
hostname router2
ip routing
int loopback0
ip address 172.16.4.1 255.255.255.0
int loopback1
ip address 172.16.5.1 255.255.255.0
int loopback 2
ip address 172.16.6.1 255.255.255.0
interface ethernet 0
ip address 172.16.10.2 255.255.255.0 no shutdown

router ospf 20
network 172.16.6.1 0.0.0.0 area 2
network 172.16.0.0 0.0.255.255 area 0
area 2 stub

```
router2#config t
Enter configuration commands, one per line.  End with CNTL/Z.
router2(config)#router ospf 20
router2(config-router)#network 172.16.0.0 0.0.255.255 area 0
router2(config-router)#network 172.16.6.1 0.0.0.0 area 2
router2(config-router)#
OSPF: "network 172.16.6.1 0.0.0.0 area 2" is ignored. It is a subset of a previous
entry.

(notice here the error.The line 172.16.0.0 0.0..255.255 area 0 places all
 172.16.*.* into area 0 prior to the line adding area 2
 You must reverse the order and put 172.16.6.1 into area 2 first !!!

router2(config-router)#no network 172.16.0.0 0.0.255.255 area 0
router2(config-router)#exit
router2(config)#
router2(config)#router ospf
% Incomplete command.

router2(config)#router ospf 20
router2(config-router)#network 172.16.6.1 0.0.0.0 area 2
router2(config-router)#network 172.16.0.0 0.0.255.255 area 0
router2(config-router)#area 2 ?
  authentication  Enable authentication
  default-cost    Set the summary default-cost of a NSSA/stub area
  nssa            Specify a NSSA area
  range           Summarize routes matching address/mask (border routers only)
  stub            Specify a stub area
  virtual-link    Define a virtual link and its parameters

router2(config-router)#area 2 stub
router2(config-router)#
router2(config-router)#^Z
router2#
%SYS-5-CONFIG_I: Configured from console by console
router2#

router2#sh run
Building configuration...

Current configuration:
!
version 11.2
no service udp-small-servers
no service tcp-small-servers
!
hostname router2
!
enable secret 5 $1$N.Tc$6HUApgc8Rc53sBdfsHe/g0
enable password 12345
!
!
interface Loopback0
 ip address 172.16.4.1 255.255.255.0
!
interface Loopback1
 ip address 172.16.5.1 255.255.255.0
!
```

```
interface Loopback2
 ip address 172.16.6.1 255.255.255.0
!
interface Ethernet0
 ip address 172.16.10.2 255.255.255.0
!
interface Serial0
 no ip address
 shutdown
!
router ospf 20
 network 172.16.0.0 0.0.0.0 area 0
 network 172.16.6.1 0.0.0.0 area 2
 network 172.16.0.0 0.0.255.255 area 0
 area 2 stub
!
no ip classless
!
line con 0
line vty 0 4
 password 12345
 login
!
end

router2#sh ip route
Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, * - candidate default
       U - per-user static route, o - ODR

Gateway of last resort is not set

     172.16.0.0/16 is variably subnetted, 7 subnets, 2 masks
C       172.16.10.0/24 is directly connected, Ethernet0
C       172.16.4.0/24 is directly connected, Loopback0
C       172.16.5.0/24 is directly connected, Loopback1
C       172.16.6.0/24 is directly connected, Loopback2
O       172.16.1.1/32 [110/11] via 172.16.10.1, 00:00:52, Ethernet0
O IA    172.16.3.1/32 [110/11] via 172.16.10.1, 00:00:52, Ethernet0
O       172.16.2.1/32 [110/11] via 172.16.10.1, 00:00:52, Ethernet0
router2#

router2#sh ip ospf

 Routing Process "ospf 20" with ID 172.16.6.1
 Supports only single TOS(TOS0) routes
 It is an area border router
 Summary Link update interval is 00:30:00 and the update due in 00:28:32
 SPF schedule delay 5 secs, Hold time between two SPFs 10 secs
 Number of DCbitless external LSA 0
 Number of DoNotAge external LSA 0
 Number of areas in this router is 2. 1 normal 1 stub 0 nssa
    Area BACKBONE(0)
       Number of interfaces in this area is 3
       Area has no authentication
       SPF algorithm executed 9 times
       Area ranges are
```

```
    Link State Update Interval is 00:30:00 and due in 00:28:32
    Link State Age Interval is 00:20:00 and due in 00:14:14
    Number of DCbitless LSA 3
    Number of indication LSA 0
    Number of DoNotAge LSA 0
  Area 2
    Number of interfaces in this area is 1
    It is a stub area
      generates stub default route with cost 1
    Area has no authentication
    SPF algorithm executed 5 times
    Area ranges are
    Link State Update Interval is 00:30:00 and due in 00:27:51
    Link State Age Interval is 00:20:00 and due in 00:15:51
    Number of DCbitless LSA 0
    Number of indication LSA 0
    Number of DoNotAge LSA 0


router2#sh ip ospf border-router

OSPF Process 20 internal Routing Table

Codes: i - Intra-area route, I - Inter-area route

i 172.16.3.1 [10] via 172.16.10.1, Ethernet0, ABR, Area 0, SPF 9
router2#
router2#sh ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
172.16.3.1        1   FULL/DR         00:00:30     172.16.10.1     Ethernet0
router2#
router2#sh ip ospf neighbor detail
 Neighbor 172.16.3.1, interface address 172.16.10.1
    In the area 0 via interface Ethernet0
    Neighbor priority is 1, State is FULL
    Options 2
    Dead timer due in 00:00:37
router2#

router2#sh ip ospf database

       OSPF Router with ID (172.16.6.1) (Process ID 20)


            Router Link States (Area 0)

Link ID         ADV Router      Age         Seq#       Checksum Link count
172.16.3.1      172.16.3.1      151         0x80000028 0x353B   3
172.16.6.1      172.16.6.1      148         0x80000030 0x8DAD   3

            Net Link States (Area 0)

Link ID         ADV Router      Age         Seq#       Checksum
172.16.10.1     172.16.3.1      151         0x80000003 0xB892

            Summary Net Link States (Area 0)

Link ID         ADV Router      Age         Seq#       Checksum
172.16.3.1      172.16.3.1      921         0x80000001 0x17C3
```

```
172.16.6.1      172.16.6.1      158         0x80000001 0xFEB5

            Router Link States (Area 2)

Link ID         ADV Router      Age         Seq#       Checksum Link count
172.16.6.1      172.16.6.1      158         0x80000004 0xDC06   1

            Summary Net Link States (Area 2)

Link ID         ADV Router      Age         Seq#       Checksum
0.0.0.0         172.16.6.1      149         0x80000001 0xEE8B
172.16.1.1      172.16.6.1      139         0x80000001 0xB8F8
172.16.2.1      172.16.6.1      139         0x80000001 0xAD03
172.16.3.1      172.16.6.1      139         0x80000001 0xA20D
172.16.4.1      172.16.6.1      144         0x80000002 0x3186
172.16.5.1      172.16.6.1      144         0x80000002 0x2690
172.16.10.0     172.16.6.1      139         0x80000001 0x5555
router2#

router2#sh ip ospf database router

       OSPF Router with ID (172.16.6.1) (Process ID 20)


            Router Link States (Area 0)

  Routing Bit Set on this LSA
  LS age: 164
  Options: (No TOS-capability, No DC)
  LS Type: Router Links
  Link State ID: 172.16.3.1
  Advertising Router: 172.16.3.1
  LS Seq Number: 80000028
  Checksum: 0x353B
  Length: 60
  Area Border Router
   Number of Links: 3

    Link connected to: a Stub Network
     (Link ID) Network/subnet number: 172.16.2.1
     (Link Data) Network Mask: 255.255.255.255
      Number of TOS metrics: 0
       TOS 0 Metrics: 1

    Link connected to: a Stub Network
     (Link ID) Network/subnet number: 172.16.1.1
     (Link Data) Network Mask: 255.255.255.255
      Number of TOS metrics: 0
       TOS 0 Metrics: 1

    Link connected to: a Transit Network
     (Link ID) Designated Router address: 172.16.10.1
     (Link Data) Router Interface address: 172.16.10.1
      Number of TOS metrics: 0
       TOS 0 Metrics: 10


  LS age: 163
  Options: (No TOS-capability, DC)
```

```
  LS Type: Router Links
  Link State ID: 172.16.6.1
  Advertising Router: 172.16.6.1
  LS Seq Number: 80000030
  Checksum: 0x8DAD
  Length: 60
  Area Border Router
   Number of Links: 3

    Link connected to: a Stub Network
     (Link ID) Network/subnet number: 172.16.5.1
     (Link Data) Network Mask: 255.255.255.255
      Number of TOS metrics: 0
       TOS 0 Metrics: 1


 Link connected to: a Stub Network
     (Link ID) Network/subnet number: 172.16.4.1
     (Link Data) Network Mask: 255.255.255.255
      Number of TOS metrics: 0
       TOS 0 Metrics: 1

    Link connected to: a Transit Network
     (Link ID) Designated Router address: 172.16.10.1
     (Link Data) Router Interface address: 172.16.10.2
      Number of TOS metrics: 0
       TOS 0 Metrics: 10



            Router Link States (Area 2)

  LS age: 175
  Options: (No TOS-capability, DC)
  LS Type: Router Links
  Link State ID: 172.16.6.1
  Advertising Router: 172.16.6.1
  LS Seq Number: 80000004
  Checksum: 0xDC06
  Length: 36
  Area Border Router
   Number of Links: 1

    Link connected to: a Stub Network
     (Link ID) Network/subnet number: 172.16.6.1
     (Link Data) Network Mask: 255.255.255.255
      Number of TOS metrics: 0
       TOS 0 Metrics: 1

router2#

router2#sh ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
172.16.3.1        1   FULL/DR         00:00:37     172.16.10.1     Ethernet0
router2#
router2#debug ip ospf adj
OSPF adjacency events debugging is on
```

(Now go & reboot ROUTE 1 so you can see the OSPF negotiation process take place)

```
router2#
OSPF: Neighbor 172.16.3.1 is dead
OSPF: neighbor 172.16.3.1 is dead, state DOWN
OSPF: Neighbor change Event on interface Ethernet0
OSPF: DR/BDR election on Ethernet0
OSPF: Elect BDR 172.16.6.1
OSPF: Elect DR 172.16.6.1
OSPF: Elect BDR 0.0.0.0
OSPF: Elect DR 172.16.6.1
       DR: 172.16.6.1 (Id)   BDR: none
OSPF: Build network LSA, router ID 172.16.6.1
OSPF: No full nbrs to build Net Lsa
OSPF: Flush network LSA on Ethernet0 for area 172.16.6.1
OSPF: Build router LSA, router ID 172.16.6.1
OSPF: Reset old DR on Ethernet0
OSPF: Build router LSA, router ID 172.16.6.1
OSPF: Build network LSA, router ID 172.16.6.1
OSPF: No full nbrs to build Net Lsa
OSPF: Flush network LSA on Ethernet0 for area 172.16.6.1
OSPF: Build router LSA, router ID 172.16.6.1
router2#
router2#
router2#
OSPF: Receive dbd from 172.16.3.1 seq 0x143D
OSPF: 2 Way Communication to neighbor 172.16.3.1
OSPF: Neighbor change Event on interface Ethernet0
OSPF: DR/BDR election on Ethernet0
OSPF: Elect BDR 172.16.3.1
OSPF: Elect DR 172.16.6.1
       DR: 172.16.6.1 (Id)   BDR: 172.16.3.1 (Id)
OSPF: send DBD packet to 172.16.10.1 seq 0x1C0A
OSPF: Build router LSA, router ID 172.16.6.1
OSPF: Receive dbd from 172.16.3.1 seq 0x1C0A
OSPF: NBR Negotiation Done  We are the MASTER
OSPF: send DBD packet to 172.16.10.1 seq 0x1C0B
OSPF: Database request to 172.16.3.1
OSPF: sent LS REQ packet to 172.16.10.1, length 12
OSPF: Receive dbd from 172.16.3.1 seq 0x1C0B
OSPF: send DBD packet to 172.16.10.1 seq 0x1C0C
OSPF: Receive dbd from 172.16.3.1 seq 0x1C0C
OSPF: Exchange Done with neighbor 172.16.3.1
OSPF: Synchronized with neighbor 172.16.3.1, state:FULL
OSPF: Build router LSA, router ID 172.16.6.1
OSPF: Build network LSA, router ID 172.16.6.1
OSPF: Neighbor change Event on interface Ethernet0
OSPF: DR/BDR election on Ethernet0
OSPF: Elect BDR 172.16.3.1
OSPF: Elect DR 172.16.6.1
       DR: 172.16.6.1 (Id)   BDR: 172.16.3.1 (Id)
OSPF: Build router LSA, router ID 172.16.6.1
OSPF: Build router LSA, router ID 172.16.6.1
router2#
router2#
```

Some OSPF terms and definitions :

Router ID: A 32 bit number used to identify each router in an ospf process. This number is typically the highest ip address assigned to each router, including both physical and loopback addresses.

Neighbor Router: A router with an interface to a common network area

Link or Interface: A connection between a network and a router interface using a single ip address/mask.

DR-Designated Router: The router designated to flood link-state information to the common network area. DR is negotiated at startup and is chosen by the highest router priority or the highest router ID.

BDR-Backup Designated Router: Router will take over if DR fails or is unavailable

Adjacency: An adjacency is formed between the DR and the BDR as well as other non-designated routers in the network area. OSPF updates are sent to routers who have formed adjacencies

Flooding: Used to synchronize and distribute link state databases between routers.

OSPF Router Types

Internal : All interfaces are defined in the same area. All internal routers have an identical link-state database in a given area.

Backbone : A backbone router has at least one interface assigned to area 0. Area 0 is always defined as the backbone area.

Area Border Router ABR : Router interfaces are connected to more than one area, or multiple area's. Information specific to each area is stored in ABR routers. An ABR must have an interface defined in each area that it is connected to.

Autonomous System Border Router ASBR : An ABR that has interfaces that connect any of the following network topologies :

1) A router that connects an internal autonomous system to some external network such as an ISP or the Internet.
2) A router that connects two or more different autonomous systems.
3) A router that connects two more different routing protocol domains, such as router that is connected to an OSPF area and also a network running RIP.
   In this case the router must have an interface in each routing domain and must be running both routing protocols. In this situation, you must redistribute or inject the other routing protocol updates into each of the routing domains.

OSPF Area Types

Backbone area : This is always defined as area 0. Backbone area's accept **all** LSA's Link State Advertisements and is used to connect multiple area's. All non-backbone area's must be physically connected to the backbone area or make use of virtual links.

Stub area : Stub area's will accept summary LSA's , but will not accept any external routing updates. Use the command **area 4 stub**

Totally Stub area : Totally stub area's will not accept summary or external updates. They contain only intra-area routes. Use the command **area 4 stub no-summary**

Standard area : normal area that accepts internal and external LSA's and summary information.

DR and BDR router election in an OSPF network

Routers in an OSPF network have an **ip ospf priority**, which by **default is set to 1** for all routers. If you set the ip ospf priority to **0**, a router will never become a DR or BDR. You can raise the ip ospf priority in order for a router to win the DR-BDR election process.
If the router priorities are the same, then the router with the highest router ID (ip address) will become the designated router. The router ID includes both the physical and loopback addresses.

Consider the following scenario:

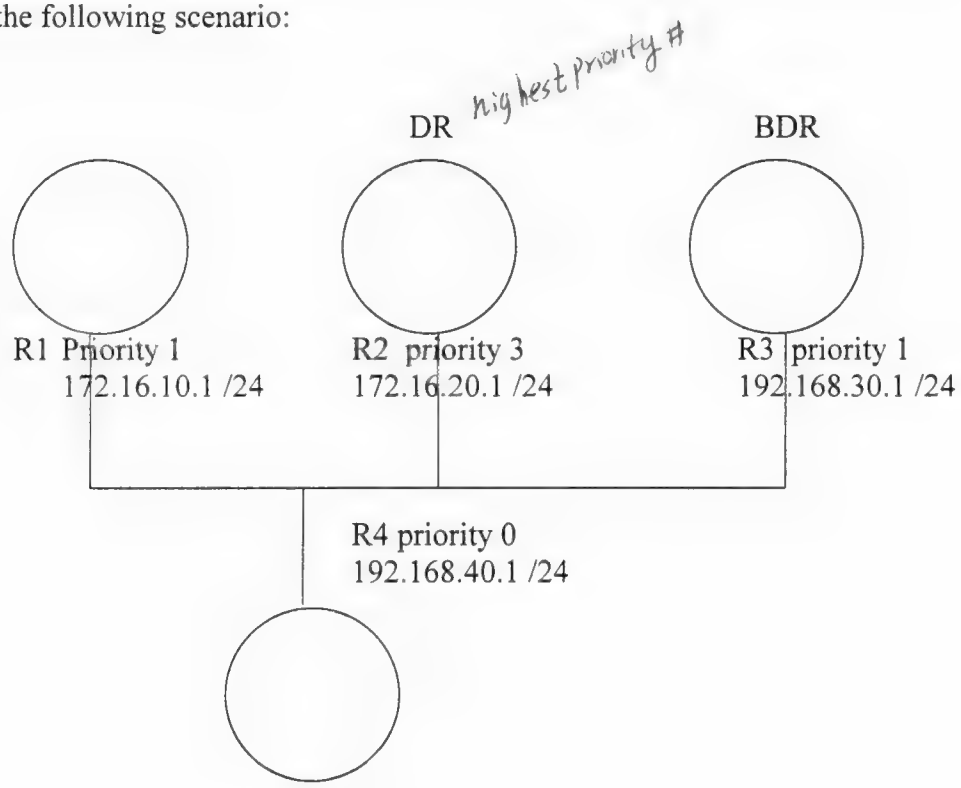

R2 wins DR because it has the highest router ID
R4 will never be a DR or BDR because it has priority 0
R3 becomes the BDR because it has a higher router ID than R1 with same priority

Using the **earlier OSPF example for multiple area's** we can analyze the DR-BDR configuration and change the ip ospf priority to switch roles.

```
router1#sh ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
172.16.6.1        1   FULL/DR         0:00:31     172.16.10.2     Ethernet0Serial0 is
administratively down, line protocol is down

router2#sh ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
172.16.3.1        1   FULL/BDR        00:00:30     172.16.10.1     Ethernet0
```

```
router1#sh ip ospf interface

Ethernet0 is up, line protocol is up
  Internet Address 172.16.10.1 255.255.255.0, Area 0
  Process ID 20, Router ID 172.16.3.1, Network Type BROADCAST, Cost: 10
  Transmit Delay is 1 sec, State BDR, Priority 1
  Designated Router (ID) 172.16.6.1, Interface address 172.16.10.2
  Backup Designated router (ID) 172.16.3.1, Interface address 172.16.10.1
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello due in 0:00:05
  Neighbor Count is 1, Adjacent neighbor count is 1
    Adjacent with neighbor 172.16.6.1  (Designated Router)
Ethernet1 is administratively down, line protocol is down
   OSPF not enabled on this interface
Loopback0 is up, line protocol is up
  Internet Address 172.16.1.1 255.255.255.0, Area 0
  Process ID 20, Router ID 172.16.3.1, Network Type LOOPBACK, Cost: 1
  Loopback interface is treated as a stub Host
Loopback1 is up, line protocol is up
  Internet Address 172.16.2.1 255.255.255.0, Area 0
  Process ID 20, Router ID 172.16.3.1, Network Type LOOPBACK, Cost: 1
  Loopback interface is treated as a stub Host
Loopback2 is up, line protocol is up
  Internet Address 172.16.3.1 255.255.255.0, Area 1
  Process ID 20, Router ID 172.16.3.1, Network Type LOOPBACK, Cost: 1
  Loopback interface is treated as a stub Host
Serial0 is administratively down, line protocol is down
   OSPF not enabled on this interface
Serial1 is administratively down, line protocol is down
   OSPF not enabled on this interface
router1#

router2#sh ip ospf interface

Ethernet0 is up, line protocol is up
  Internet Address 172.16.10.2/24, Area 0

  Process ID 20, Router ID 172.16.6.1, Network Type BROADCAST, Cost: 10
  Transmit Delay is 1 sec, State DR, Priority 1
  Designated Router (ID) 172.16.6.1, Interface address 172.16.10.2
  Backup Designated router (ID) 172.16.3.1, Interface address 172.16.10.1
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello due in 00:00:06
  Neighbor Count is 1, Adjacent neighbor count is 1
    Adjacent with neighbor 172.16.3.1  (Backup Designated Router)
  Suppress hello for 0 neighbor(s)
Loopback0 is up, line protocol is up
  Internet Address 172.16.4.1/24, Area 0
  Process ID 20, Router ID 172.16.6.1, Network Type LOOPBACK, Cost: 1
  Loopback interface is treated as a stub Host
Loopback1 is up, line protocol is up
  Internet Address 172.16.5.1/24, Area 0
  Process ID 20, Router ID 172.16.6.1, Network Type LOOPBACK, Cost: 1
  Loopback interface is treated as a stub Host
Loopback2 is up, line protocol is up
  Internet Address 172.16.6.1/24, Area 2
  Process ID 20, Router ID 172.16.6.1, Network Type LOOPBACK, Cost: 1
  Loopback interface is treated as a stub Host
```

```
Serial0 is administratively down, line protocol is down
   OSPF not enabled on this interface
router2#
```

(Now we are going to raise the current BDR router priority (router1) so that it will become the new DR)

```
router1#config t

Enter configuration commands, one per line.  End with CNTL/Z.
router1(config)#int e0
router1(config-if)#ip ospf ?
  authentication-key   Authentication password (key)
  cost                 Interface cost
  dead-interval        Interval after which a neighbor is declared dead
  hello-interval       Time between HELLO packets
  message-digest-key   Message digest authentication password (key)
  network              Network type
  priority             Router priority
  retransmit-interval  Time between retransmitting lost link state
                       advertisements
  transmit-delay       Link state transmit delay

router1(config-if)#ip ospf priority ?
  <0-255>  Priority

router1(config-if)#ip ospf priority 5
router1(config-if)#^Z
router1#
%SYS-5-CONFIG_I: Configured from console by console
router1#
```

Here you see that Router1 is now the DR because of the higher ip ospf priority

```
router1#sh ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
172.16.6.1        1   FULL/BDR        0:00:31     172.16.10.2     Ethernet0
router1#
router2#sh ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
172.16.3.1        5   FULL/DR         00:00:37     172.16.10.1     Ethernet0

router2#
```

router1# sh ip ospf interface

```
Ethernet0 is up, line protocol is up
  Internet Address 172.16.10.1 255.255.255.0, Area 0
  Process ID 20, Router ID 172.16.3.1, Network Type BROADCAST, Cost: 10
  Transmit Delay is 1 sec, State DR, Priority 5
  Designated Router (ID) 172.16.3.1, Interface address 172.16.10.1
  Backup Designated router (ID) 172.16.6.1, Interface address 172.16.10.2
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello due in 0:00:07
  Neighbor Count is 1, Adjacent neighbor count is 1
    Adjacent with neighbor 172.16.6.1  (Backup Designated Router)
Ethernet1 is administratively down, line protocol is down
   OSPF not enabled on this interface
```

```
Loopback0 is up, line protocol is up
  Internet Address 172.16.1.1 255.255.255.0, Area 0
  Process ID 20, Router ID 172.16.3.1, Network Type LOOPBACK, Cost: 1
  Loopback interface is treated as a stub Host
Loopback1 is up, line protocol is up
  Internet Address 172.16.2.1 255.255.255.0, Area 0
  Process ID 20, Router ID 172.16.3.1, Network Type LOOPBACK, Cost: 1
  Loopback interface is treated as a stub Host
Loopback2 is up, line protocol is up
  Internet Address 172.16.3.1 255.255.255.0, Area 1
  Process ID 20, Router ID 172.16.3.1, Network Type LOOPBACK, Cost: 1
  Loopback interface is treated as a stub Host
Serial0 is administratively down, line protocol is down
   OSPF not enabled on this interface
Serial1 is administratively down, line protocol is down
   OSPF not enabled on this interface
router1#
```

```
router2# sh ip ospf interface
Ethernet0 is up, line protocol is up
  Internet Address 172.16.10.2/24, Area 0
  Process ID 20, Router ID 172.16.6.1, Network Type BROADCAST, Cost: 10
  Transmit Delay is 1 sec, State BDR, Priority 1
  Designated Router (ID) 172.16.3.1, Interface address 172.16.10.1
  Backup Designated router (ID) 172.16.6.1, Interface address 172.16.10.2
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello due in 00:00:00
  Neighbor Count is 1, Adjacent neighbor count is 1
    Adjacent with neighbor 172.16.3.1  (Designated Router)
  Suppress hello for 0 neighbor(s)
Loopback0 is up, line protocol is up
  Internet Address 172.16.4.1/24, Area 0
  Process ID 20, Router ID 172.16.6.1, Network Type LOOPBACK, Cost: 1
  Loopback interface is treated as a stub Host
Loopback1 is up, line protocol is up
  Internet Address 172.16.5.1/24, Area 0
  Process ID 20, Router ID 172.16.6.1, Network Type LOOPBACK, Cost: 1
  Loopback interface is treated as a stub Host
Loopback2 is up, line protocol is up
  Internet Address 172.16.6.1/24, Area 2
  Process ID 20, Router ID 172.16.6.1, Network Type LOOPBACK, Cost: 1
  Loopback interface is treated as a stub Host
Serial0 is administratively down, line protocol is down
   OSPF not enabled on this interface
router2#
```

Route summarization for multiple ospf area's

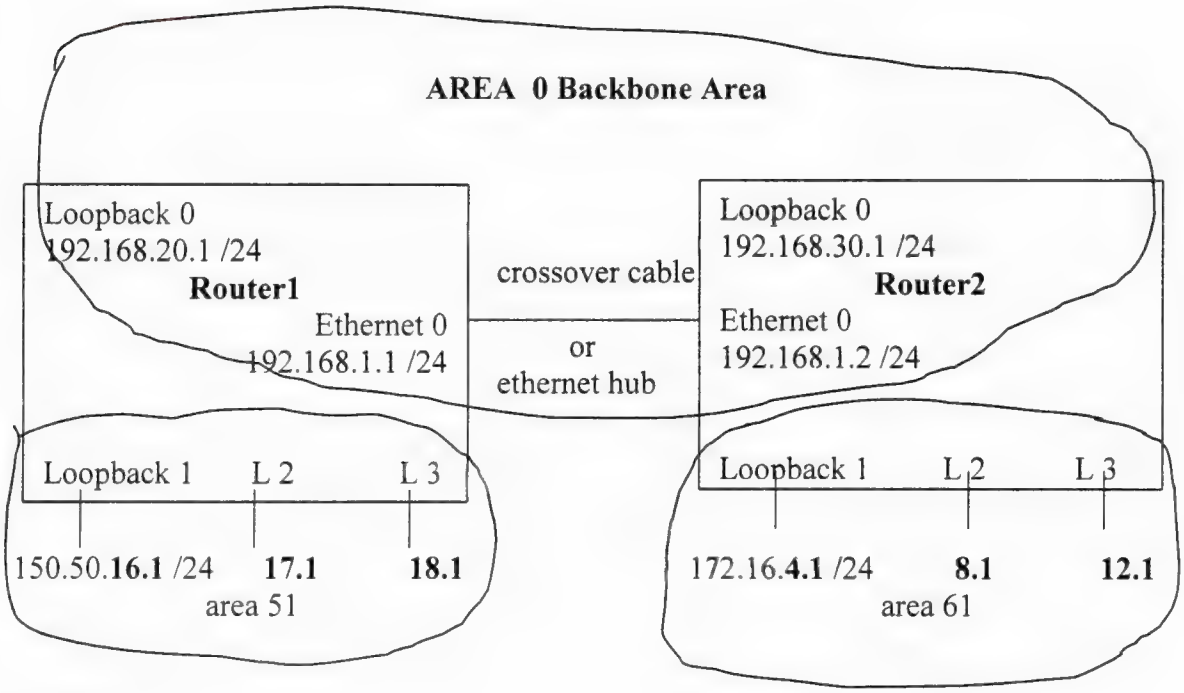

Config t
Hostname router1
Ip routing
Int loopback0
Ip address 192.168.20.1 255.255.255.0
Int loopback1
Ip address 150.50.16.1 255.255.255.0
Int loopback 2
Ip address 150.50.17.1 255.255.255.0
Int loopback 3
Ip address 150.50.18.1 255.255.255.0
int ethernet 0
Ip address 192.168.1.1 255.255.255.0

Router ospf 30
Network 192.168.1.0 0.0.0.255 area 0
Network 192.168.20.1 0.0.0.0 area 0
Network 150.50.16.1 0.0.0.0 area 51
Network 150.50.17.1 0.0.0.0 area 51
Network 150.50.18.1 0.0.0.0 area 51

Config t
hostname router2
ip routing
int loopback0
ip address 192.168.30.1 255.255.255.0
int loopback1
ip address 172.16.4.1 255.255.255.0
int loopback 2
ip address 172.16.8.1 255.255.255.0
int loopback 3
ip address 172.16.12.1 255.255.255.0 int ethernet 0
ip address 192.168.1.2 255.255.255.0

router ospf 30
network 192.168.1.0 0.0.0.255 area 0
network 192.168.30.1 0.0.0.0 area 0
network 172.16.4.1 0.0.0.0 area 61
network 172.16.8.1 0.0.0.0 area 61
network 172.16.12.1 0.0.0.0 area 61

```
router1#show running-config
Building configuration...

Current configuration:
!!
hostname router1
!
enable secret 5 $1$rGRF$7Ve8JHSHDyVnXob6EilcV1
enable password 12345
!
!
interface Loopback0
 ip address 192.168.20.1 255.255.255.0
!
interface Loopback1
 ip address 150.50.16.1 255.255.255.0
!
interface Loopback2
 ip address 150.50.17.1 255.255.255.0
!
interface Loopback3
 ip address 150.50.18.1 255.255.255.0
!
interface Ethernet0
 ip address 192.168.1.1 255.255.255.0
!
interface Ethernet1
 no ip address
 shutdown
!
interface Serial0
 no ip address
 shutdown
 no fair-queue
!
interface Serial1
 no ip address
 shutdown
!
router ospf 30
 network 150.50.16.1 0.0.0.0 area 51
 network 150.50.17.1 0.0.0.0 area 51
 network 150.50.18.1 0.0.0.0 area 51
 network 192.168.1.0 0.0.0.255 area 0
 network 192.168.20.1 0.0.0.0 area 0
!
line con 0
line aux 0
 transport input all
line vty 0 4
 password 12345
 login
!
end

router1#
```

```
router2#
router2#show run
Building configuration...

Current configuration:
hostname router2
!
enable secret 5 $1$j/nL$0rOxRjSJj.N2X2/A7OpWx0
enable password 12345
!
!
interface Loopback0
 ip address 192.168.30.1 255.255.255.0
!
interface Loopback1
 ip address 172.16.4.1 255.255.255.0
!
interface Loopback2
 ip address 172.16.8.1 255.255.255.0
!
interface Loopback3
 ip address 172.16.12.1 255.255.255.0
!
interface Ethernet0
 ip address 192.168.1.2 255.255.255.0
!
interface Serial0
 no ip address
 shutdown
!
router ospf 30
 network 172.16.4.1 0.0.0.0 area 61
 network 172.16.8.1 0.0.0.0 area 61
 network 172.16.12.1 0.0.0.0 area 61
 network 192.168.1.0 0.0.0.255 area 0
 network 192.168.30.1 0.0.0.0 area 0
!
no ip classless
!
line con 0
line vty 0 4
 password 12345
 login
!
end

router2#
```

Here you can see that both routers have learned about 4 ospf routes from the other router (1 regular ospf and 3 ospf inter area routes ***O-IA****.*

```
router1#
router1#show ip route
Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, * - candidate default

Gateway of last resort is not set

C    192.168.1.0 is directly connected, Ethernet0
     192.168.30.0 255.255.255.255 is subnetted, 1 subnets
O       192.168.30.1 [110/11] via 192.168.1.2, 00:02:44, Ethernet0
C    192.168.20.0 is directly connected, Loopback0
     150.50.0.0 255.255.255.0 is subnetted, 3 subnets
C       150.50.16.0 is directly connected, Loopback1
C       150.50.17.0 is directly connected, Loopback2
C       150.50.18.0 is directly connected, Loopback3
     172.16.0.0 255.255.255.255 is subnetted, 3 subnets
O IA    172.16.12.1 [110/11] via 192.168.1.2, 00:02:45, Ethernet0
O IA    172.16.8.1 [110/11] via 192.168.1.2, 00:02:45, Ethernet0
O IA    172.16.4.1 [110/11] via 192.168.1.2, 00:02:45, Ethernet0

router2#
router2#show ip route
Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, * - candidate default
       U - per-user static route, o - ODR

Gateway of last resort is not set

C    192.168.1.0/24 is directly connected, Ethernet0
C    192.168.30.0/24 is directly connected, Loopback0
     192.168.20.0/32 is subnetted, 1 subnets
O       192.168.20.1 [110/11] via 192.168.1.1, 00:03:11, Ethernet0
     150.50.0.0/32 is subnetted, 3 subnets
O IA    150.50.17.1 [110/11] via 192.168.1.1, 00:03:11, Ethernet0
O IA    150.50.16.1 [110/11] via 192.168.1.1, 00:03:11, Ethernet0
O IA    150.50.18.1 [110/11] via 192.168.1.1, 00:03:11, Ethernet0
     172.16.0.0/24 is subnetted, 3 subnets
C       172.16.12.0 is directly connected, Loopback3
C       172.16.8.0 is directly connected, Loopback2
C       172.16.4.0 is directly connected, Loopback1
router2#
```

Routes displayed - without summarization

A better way to do this is to summarize these 3 inter area routes into one ospf inter area route to reduce the size of the routing table. Your instructor will explain the details and theory of route summarization.

There are two ways to summarize routes:
1) For ***external*** *routes use the* ***summary-address*** *command on* ***ASBR's***
2) For ***inter-area*** *routes use the* ***area range*** *command on* ***ABR's***

We will use the area range command on the ABR's for areas 51 and 61
Instructor will explain why this works !!

```
router1#config t
Enter configuration commands, one per line.  End with CNTL/Z.
router1(config)#router ospf 30
router1(config-router)#area ?
  <0-4294967295>  OSPF area ID as a decimal value
  A.B.C.D         OSPF area ID in IP address format

router1(config-router)#area 51 ?
  authentication  Enable authentication
  default-cost    Set the default-cost of a stub area
  range           Summarize routes matching address/mask (border routers only)
  stub            Specify a stub area
  virtual-link    Define a virtual link and its parameters

router1(config-router)#area 51 range ?
  A.B.C.D  IP address to match

router1(config-router)#area 51 range 150.50.16.0 ?
  A.B.C.D  IP mask for address

router1(config-router)#area 51 range 150.50.16.0 255.255.252.0
router1(config-router)#
router1(config-router)#^Z
router1#
%SYS-5-CONFIG_I: Configured from console by console
```

/28

We summarize the routes

```
router2#
router2#config t
Enter configuration commands, one per line.  End with CNTL/Z.
router2(config)#router ospf 30
router2(config-router)#
router2(config-router)#area 61 range 172.16.0.0 255.255.240.0
router2(config-router)#^Z
router2#
%SYS-5-CONFIG_I: Configured from console by console
router2#
```

/20

We summarize the routes

255.255.1111 0000.0

summarize the bits.
common bits
shift left

Here you can see that the 3 ospf inter area routes have been summarized into 1 inter area route, thus reducing the size of the routing table.

```
router1#show ip route
Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, * - candidate default

Gateway of last resort is not set

C    192.168.1.0 is directly connected, Ethernet0
     192.168.30.0 255.255.255.255 is subnetted, 1 subnets
O       192.168.30.1 [110/11] via 192.168.1.2, 00:16:03, Ethernet0
C    192.168.20.0 is directly connected, Loopback0
     150.50.0.0 255.255.255.0 is subnetted, 3 subnets
C       150.50.16.0 is directly connected, Loopback1
C       150.50.17.0 is directly connected, Loopback2
```

```
C       150.50.18.0 is directly connected, Loopback3
     172.16.0.0 255.255.240.0 is subnetted, 1 subnets
O IA    172.16.0.0 [110/11] via 192.168.1.2, 00:07:28, Ethernet0
```

Verify that you can still ping interfaces in area 61 on router2

```
router1#
router1#ping 172.16.8.1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.8.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/3/4 ms
router1#
router1#
```

Likewise for router2

```
router2#
router2#show ip route
Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, * - candidate default
       U - per-user static route, o - ODR

Gateway of last resort is not set

C    192.168.1.0/24 is directly connected, Ethernet0
C    192.168.30.0/24 is directly connected, Loopback0
     192.168.20.0/32 is subnetted, 1 subnets
O       192.168.20.1 [110/11] via 192.168.1.1, 00:08:21, Ethernet0
     150.50.0.0/22 is subnetted, 1 subnets
O IA    150.50.16.0 [110/11] via 192.168.1.1, 00:08:21, Ethernet0
     172.16.0.0/24 is subnetted, 3 subnets
C       172.16.12.0 is directly connected, Loopback3
C       172.16.8.0 is directly connected, Loopback2
C       172.16.4.0 is directly connected, Loopback1
router2#
router2#ping 150.50.17.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 150.50.17.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/3/4 ms
router2#
router2#

router1#

router1#show run
Building configuration...

Current configuration:
!
version 11.0
service udp-small-servers
service tcp-small-servers
!
hostname router1
```

```
!
enable secret 5 $1$rGRF$7Ve8JHSHDyVnXob6EilcV1
enable password 12345
!
!
interface Loopback0
 ip address 192.168.20.1 255.255.255.0
!
interface Loopback1
 ip address 150.50.16.1 255.255.255.0
!
interface Loopback2
 ip address 150.50.17.1 255.255.255.0
!
interface Loopback3
 ip address 150.50.18.1 255.255.255.0
!
interface Ethernet0
 ip address 192.168.1.1 255.255.255.0
!
interface Ethernet1
 no ip address
 shutdown
!
interface Serial0
 no ip address
 shutdown
 no fair-queue
!
interface Serial1
 no ip address
 shutdown
!
router ospf 30
 network 150.50.16.1 0.0.0.0 area 51
 network 150.50.17.1 0.0.0.0 area 51
 network 150.50.18.1 0.0.0.0 area 51
 network 192.168.1.0 0.0.0.255 area 0
 network 192.168.20.1 0.0.0.0 area 0
 area 51 range 150.50.16.0 255.255.252.0
!
!
!
end

router2#
router2#show run
Building configuration...

Current configuration:
!
version 11.2
no service udp-small-servers
no service tcp-small-servers
!
hostname router2
!
enable secret 5 $1$j/nL$0rOxRjSJj.N2X2/A7OpWx0
enable password 12345
!
```

```
!
interface Loopback0
 ip address 192.168.30.1 255.255.255.0
!
interface Loopback1
 ip address 172.16.4.1 255.255.255.0
!
interface Loopback2
 ip address 172.16.8.1 255.255.255.0
!
interface Loopback3
 ip address 172.16.12.1 255.255.255.0
!
interface Ethernet0
 ip address 192.168.1.2 255.255.255.0
!
interface Serial0
 no ip address
 shutdown
!
router ospf 30
 network 172.16.4.1 0.0.0.0 area 61
 network 172.16.8.1 0.0.0.0 area 61
 network 172.16.12.1 0.0.0.0 area 61
 network 192.168.1.0 0.0.0.255 area 0
 network 192.168.30.1 0.0.0.0 area 0
 area 61 range 172.16.0.0 255.255.240.0
!
!
end
```

Variable Length Subnet Masking (VLSM)

As we mentioned in the section "Working with Subnet Masks", subnet masks are based on the binary number system. For example we said that:

binary 11111111 = decimal 255

The question is, how did we do this conversion? First, some termonology:

- A bit is a single binary character. Acceptable values are "1" or "0".
- A byte is a collection of 8 binary characters (or 8 bits).

This is why IP addresses are referred to as having "32 bits of address space". An IP address is simply 4 bytes seperated by periods. Doing the math, 4 x 8 bits =32 bits

So why does a binary 11111111 equal a decimal 255?

Within a binary byte, each of the bits is assigned a decimal value. The values are shown in the following diagram:

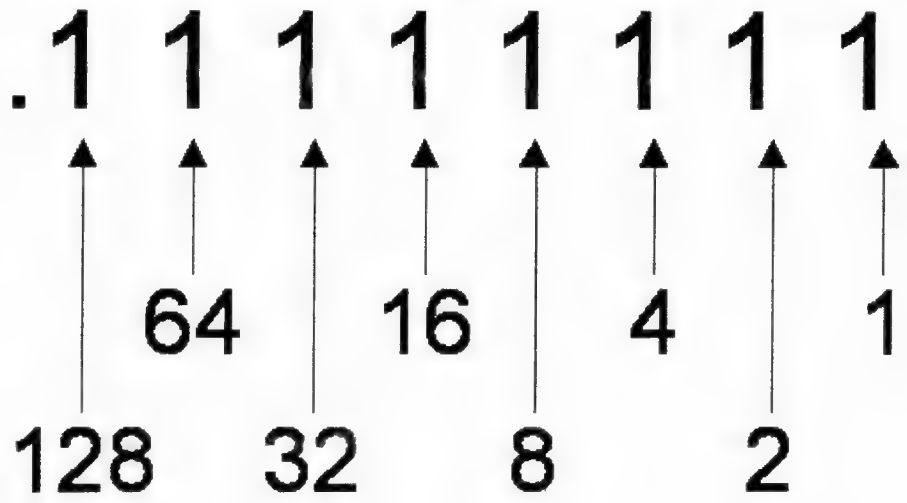

So if I see a binary value of 10000000, I know that the decimal equivelant is 128
If I see a binary value of 00000001, I know that the decimal equivelant is 1
If I see a binary value of 11111111, I know that the decimal equivelant is 255
(128+64+32+16+8+4+2+1 = 255)

When you are working with IP addresses, you will never see an address or a subnet mask that has a value higher than 255. This is because 255 is the largest number you can express in binary format.

A regular subnet mask has all bits set either high or low. Thus the only values produced are 0 and 255. When a value of 255 is used, it is assumed that the entire byte of the IP address is used for networking. If a value of 0 is used, it assumes that the corresponding byte is used to identify the host.

For example:

```
192.168.100.5
255.255.255.0
```

In this example the first three bytes of the subnet mask are all set to 255. This means that 192.168.100 identifies the network address. The IP address for all machines on this subnet must start with 192.168.100.

The last byte is set to 0. This means that "5" is the unique host identifier for the system using this IP address. No other system on the 192.168.100 network should be allowed to use a value of 5 for the host address.

In binary format, this subnet mask would look like this:
11111111.11111111.11111111.00000000

Typically, when an byte contains all zeros, only a single zero is used. This means this subnet mask can also appear as followed:
11111111.11111111.11111111.0

With a variable length subnet mask, a byte can contain a mixture of ones and zeros. In other words we "borrow" a portion of the host byte in order to create additional network space.

Consider the following example:

Network address: 200.100.50.0
decimal subnet mask: 255.255.255.0
binary subnet mask: 11111111.11111111.11111111.0

For example purposes, let's assume that we have been assigned the above address space by our ISP. Let us further assume that we actually have two separate networks that we need to assign IP addresses to. This creates a problem as we have only one network number but two physical networks.

To resolve this problem, we need to borrow a bit from the host portion of the address in order to split this network address space in half.

So we start with:

Network address: 200.100.50.0
decimal subnet mask: 255.255.255.0
binary subnet mask: 11111111.11111111.11111111.0

And end up with:

Network address: 200.100.50.0
decimal subnet mask: 255.255.255.128
binary subnet mask: 11111111.11111111.11111111.10000000

Network address: 200.100.50.128
decimal subnet mask: 255.255.255.128
binary subnet mask: 11111111.11111111.11111111.10000000

Valid hosts on our first network would be:
200.100.50.1 - 200.100.50.126

Valid hosts on our second network would be:
200.100.50.129 - 200.100.50.254

Where did all these numbers come from?. Let's take it one step at a time to see how we arrived at all these values.

First, we change the binary version of the subnet mask from:
11111111.11111111.11111111.0

To:
11111111.11111111.11111111.10000000

This was done so that more of the address space could be used for creating networks. Of course this also means that less of the address space is available for identifying hosts.

We have gone from:

```
11111111.11111111.11111111.00000000
|<--  Network portion  -->|<-host->|

To:
11111111.11111111.11111111.10000000
|<--   Network portion   -->| host |
```

Earlier we mentioned that the first bit in a byte is equivalent to a decimal 128. This is why the decimal equivelant of the subnet mask changed from:
255.255.255.0

To:
255.255.255.128

Since the first bit in the last byte is now part of the network address space, setting this bit low or high allows us to identify two separate network.

For example we started with:
200.100.50.0

In binary, the address would look like this:
11001000.1100100.110010.0

This remains the network address for our first network because the first bit is set low.

If we trigger the first bit high however we get:
200.100.50.128

This is because in binary the address looks like this:
11001000.1100100.110010.10000000

Since we have "borrowed" one bit in the last byte for the network portion of the address, it stands to reason that we now have one less bit to work with when addressing each of our hosts. Returning to our figure:

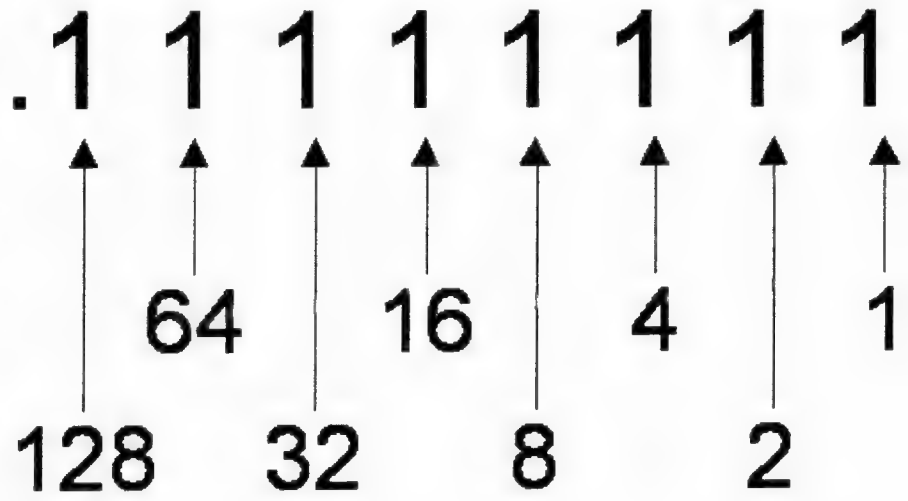

The bits with a decimal equivelant of 1-64 can still be used to identify hosts. Adding these values we get:
1+2+4+8+16+32+64=127

So we have 127 values free to work with when addressing each of our new subnets. This means that we can have 126 hosts on each of our two networks.

Why 126 hosts instead of 127?

We can only have 126 hosts on each subnet because in IP the last host value is always used as a broadcast address.

Now the math wizzes out there are probably saying:

"Wait a minute. I've lost some address space. With a full class C I can have 254 hosts, but now that we've subnetted I only have 2x126=252 hosts. What gives?"

This is 100% correct. Remember that every network has to have a network address and a broadcast address. When every you use VLSM to create additional networks, you are going to lose numbers numbers to fill these roles.

So in our example:
200.100.50.127 - Can not be used to identify a host as it has become the broadcast address for the first subnet.
200.100.50.128 - Can not be used to identify a host as it has become the network address for the second subnet.

The more bits you borrow for network address space, the more host numbers you will lose for identifying network and broadcast addresses.

The following table shows how a class C subnet would be broken up using different mask values. When more than four subnets have been created, we have only shown address space for the first and the last two subnets.

Mask (last byte in binary)	# of Subnets	Hosts per Subnet	Available host Address ranges	Network Address value	Broadcast Address value
255.255.255.128 (.10000000)	2	126	1 – 126 129 – 254	0 128	127 255
255.255.255.192 (.11000000)	4	62	1 – 62 65 – 126 129 – 190 193 – 254	0 64 128 192	63 127 191 255
255.255.255.224 (.11100000)	8	30	1 – 30 33 – 62 193 – 222 225 – 254	0 32 192 224	31 63 223 255
255.255.255.240 (.11110000)	16	14	1 – 14 17 – 30 225 – 238 241 – 254	0 16 224 240	15 31 239 255
255.255.255.248 (.11111000)	32	6	1 – 6 9 – 14 241 – 246 249 – 254	0 8 240 248	7 15 247 255
255.255.255.252 (.11111100)	64	2	1 – 2 5 – 6 249 – 250 253 – 254	0 4 248 252	3 7 251 255
255.255.255.254 (.11111110)	128	0	0	0	0

- It is worth noting that there are some discrepancies regarding whether it is okay to use "subnet zero" when using VLSM
- Subnet zero refers to the first and last subnet created when a non-standard mask is applied

- Per RFC 1700, the first and last network created when VLSM is used is suppose to be dropped (not used)

This means that our table from the last page would be changed as followed:

Mask	# of Subnets
255.255.255.128	0
255.255.255.192	2
255.255.255.224	6
255.255.255.240	14
255.255.255.248	30
255.255.255.252	62

- Dropping subnet zero means that "legally", a 255.255.255.128 subnet mask can not be used as both subnets created must be dropped
- This configuration was required due to routing limitation in older UNIX machines
- Some older operating systems will not communicate with, or accept, a zero subnet address (NetWare 3.1x)
- A Cisco router must have zero subnetting specifically enabled in order to accept these addresses
- All modern day operating systems will accept zero subnet addresses
- Despite that fact that modern systems will support subnet zero, many vendor certification course still insist it not be used
- Recent RFC's have declared that it is acceptable to use subnet zero when addressing networks

Network administrators will typically avoid using subnet zero when possible, but deploy it when address space limitations require that it be used.

A 255.255.255.128 Subnetted Network

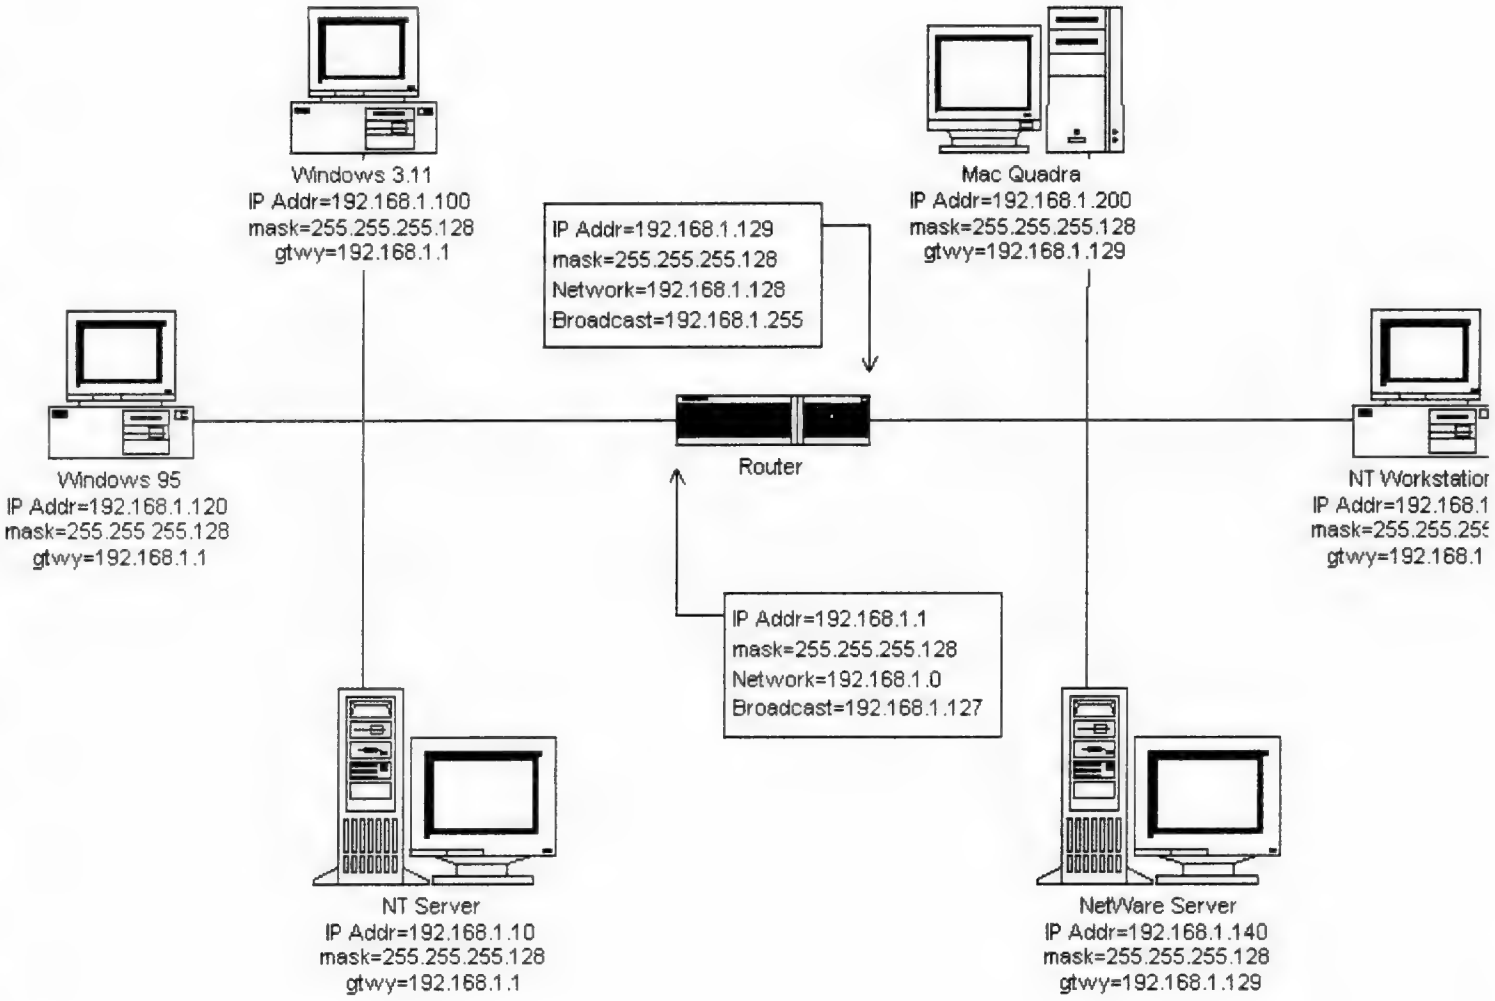

Advanced VLSM

- The use of variable length subnet masks can be mixed and matched
- You are not required to use a consistent mask value on every subnet on your network
- For example you could break up a legal class C by using /27 on your LAN segments and /30 on your WAN links
- The Routing Information Protocol (RIP) does not support VLSM as it expects all networks to use a consistent mask value
- Lets look at an example of how you can mix subnet values.

  Let's assume that your ISP has issued you a legal address space using a 255.255.255.192 mask (/26)

- You wish to install a firewall which will create an additional subnet, but your ISP will not allocate you additional address space

- You can not use Network Address Translation (NAT) because you need to support a number of DCOM applications such as NetMeeting and Exchange
- You only have 20 internal hosts which require an IP address

A drawing of the network setup would look something like this:

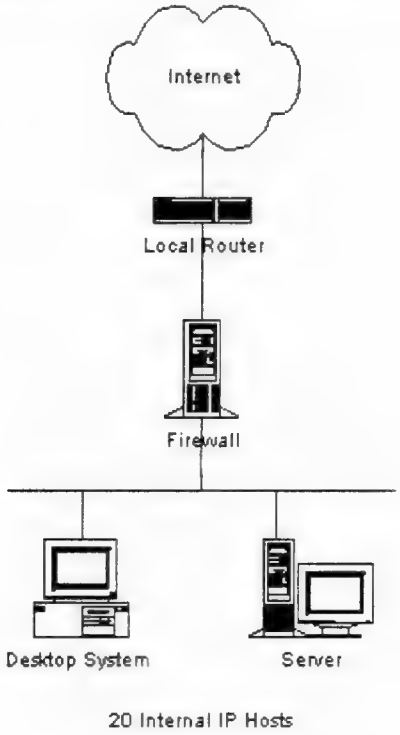

- So we need to break up the address space that has been allocated to us in a way that will allow us to address all our network segments.
- For the purpose of example, assume the address space given to us by our ISP is 192.168.1.64/26

This address space would look like this:
Network address: 192.168.1.64
Subnet mask: 255.255.255.192 (/26)
Broadcast address: 192.168.1.127
Valid Hosts: 192.168.1.65 - 192.168.1.126
We would need to break up the 192.168.1.64/26 address space as followed:

- First, we adjust the mask value to the smallest subnet size that will support our 20 internal hosts
- This would be a mask value of 255.255.255.224 (/27) which will allow us to address 30 systems

So our first subnet would be as followed:
Network address: 192.168.1.64
Subnet mask: 255.255.255.224 (/27)
Broadcast Address: 192.168.1.95
Valid hosts: 192.168.1.65 - 192.168.1.94

- Remember that the original mask value gave us addresses up to 192.168.1.127 to work with
- This means we still have 192.168.1.95 - 192.168.1.127 to work with
- We also need to address the subnet between the firewall and the router
- This network segment only needs to support two IP addresses

Jumping to the end of the available address space, our second subnet would look something like this:
Network address: 192.168.1.124
Subnet mask: 255.255.255.252 (/30)
Broadcast Address: 192.168.1.127
Valid hosts: 192.168.1.125 - 192.168.1.126

This leaves the address space 192.168.96 - 192.168.1.123 unused

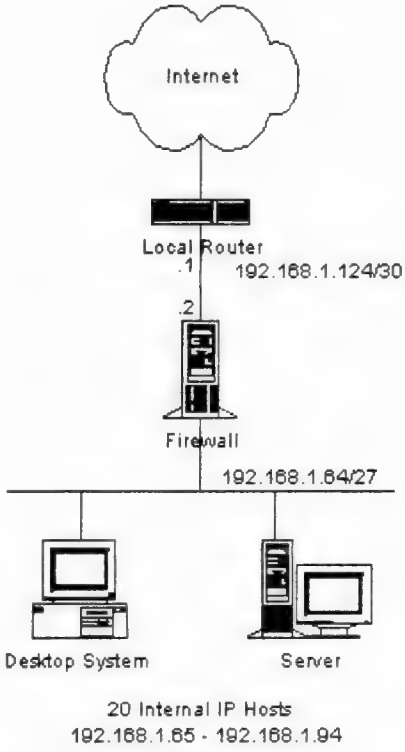

So what about the address space 192.168.96 - 192.168.1.123? Does this go to waste?
Unfortunately, yes unless you later create additional subnets.

For example, what if you later decide to add a third NIC to the firewall to create a DMZ?

Using the remaining address space, you have a couple of options:
192.168.1.96/28 - This would support 14 hosts
192.168.1.96/29 - This would support 6 hosts, but leave address space free for more additional subnets

Variable Length Subnet masking can be a great way to insure that you get the most effective use available out of your address space.

```
router1#show queueing
Current fair queue configuration:
interface Serial0

  Input queue: 0/75/0 (size/max/drops); Total output drops: 0
  Output queue: 0/96/0 (size/threshold/drops)
     Conversations  0/1 (active/max active)
     Reserved Conversations 0/0 (allocated/max allocated)

Current priority queue configuration:
Current custom queue configuration:


router1#show queueing fair
Current fair queue configuration:
interface Serial0

  Input queue: 0/75/0 (size/max/drops); Total output drops: 0
  Output queue: 0/96/0 (size/threshold/drops)
     Conversations  0/1 (active/max active)
     Reserved Conversations 0/0 (allocated/max allocated)


router1#show queue ?
  Ethernet  IEEE 802.3
  Loopback  Loopback interface
  Null      Null interface
  Serial    Serial

router1#show queue serial 0
  Input queue: 0/75/0 (size/max/drops); Total output drops: 0
  Output queue: 0/96/0 (size/threshold/drops)
     Conversations  0/1 (active/max active)
     Reserved Conversations 0/0 (allocated/max allocated)

router1#show interface serial 0
Serial0 is up, line protocol is up
  Hardware is HD64570
  Internet address is 172.16.3.1 255.255.255.0
  MTU 1500 bytes, BW 1544 Kbit, DLY 20000 usec, rely 255/255, load 1/255
  Encapsulation HDLC, loopback not set, keepalive set (10 sec)
  Last input 0:00:00, output 0:00:02, output hang never
  Last clearing of "show interface" counters never
  Input queue: 0/75/0 (size/max/drops); Total output drops: 0
  Output queue: 0/96/0 (size/threshold/drops)
     Conversations  0/1 (active/max active)
     Reserved Conversations 0/0 (allocated/max allocated)
  5 minute input rate 0 bits/sec, 0 packets/sec
  5 minute output rate 0 bits/sec, 0 packets/sec
     5308 packets input, 336765 bytes, 0 no buffer
```

Configuring Priority queuing

```
router2#show queueing
Current fair queue configuration:

  Interface     Discard       Dynamic        Reserved
                threshold     queue count    queue count
  Serial0             64             256              0

Current priority queue configuration:
Current custom queue configuration:
Current RED queue configuration:
router2#
router2#config t
Enter configuration commands, one per line.  End with CNTL/Z.
router2(config)#interface serial 0
router2(config-if)#no fair-queue
router2(config-if)#^Z
router2#
router2#show queueing
Current fair queue configuration:
Current priority queue configuration:
Current custom queue configuration:
Current RED queue configuration:

router2#
router2#config t
Enter configuration commands, one per line.  End with CNTL/Z.
router2(config)#
router2(config)#priority-list 1 ?
  default      Set priority queue for unspecified datagrams
  interface    Establish priorities for packets from a named interface
  protocol     priority queueing by protocol
  queue-limit  Set queue limits for priority queues

router2(config)#priority-list 1 protocol ?
  arp            IP ARP
  bridge         Bridging
  cdp            Cisco Discovery Protocol
  compressedtcp  Compressed TCP
  ip             IP
  llc2           llc2
  pad            PAD links
  snapshot       Snapshot routing support


router2(config)#priority-list 1 protocol ip ?
  high
  medium
  normal
  low



router2(config)#priority-list 1 protocol ip high ?
  fragments  Prioritize fragmented IP packets
  gt         Prioritize packets greater than a specified size
  list       To specify an access list
  lt         Prioritize packets less than a specified size
  tcp        Prioritize TCP packets 'to' or 'from' the specified port
```

```
  udp        Prioritize UDP packets 'to' or 'from' the specified port
  <cr>

router2(config)#priority-list 1 protocol ip high gt ?
  <0-65535>  Packet size (include MAC encapsulation bytes)

router2(config)#priority-list 1 protocol ip high gt 1500
router2(config)#priority-list 1 protocol ip low lt 256
router2(config)#priority-list 1 protocol ip normal
router2(config)#
router2(config)#priority-list 1 interface ?
  Ethernet  IEEE 802.3
  Loopback  Loopback interface
  Null      Null interface
  Serial    Serial

router2(config)#priority-list 1 interface serial 0 ?
  high
  medium
  normal
  low

router2(config)#priority-list 1 interface serial 0 normal
router2(config)#priority-list 1 interface ethernet 0 medium
router2(config)#priority-list 1 default normal

router2(config)#priority-list 1 queue-limit ?
  <0-32767>  High limit

router2(config)#priority-list 1 queue-limit 40 ?
  <0-32767>  Medium limit

router2(config)#priority-list 1 queue-limit 40 80 ?
  <0-32767>  Normal limit

router2(config)#priority-list 1 queue-limit 40 80 120 ?
  <0-32767>  Lower limit

router2(config)#priority-list 1 queue-limit 40 80 120 160
router2(config)#
router2(config)#interface serial 0
router2(config-if)#priority-group 1
router2(config-if)#
router2(config-if)#^Z

router2#show queueing ?
  custom    custom queueing list configuration
  fair      fair queueing configuration
  priority  priority queueing list configuration
  red       random early detection configuration
  <cr>

router2#show queueing priority
Current priority queue configuration:

List   Queue  Args
1      high   protocol ip          gt 1500
1      low    protocol ip          lt 256
1      normal protocol ip
1      normal interface Serial0
1      medium interface Ethernet0
```

```
1       high   limit 40
1       medium limit 80
1       normal limit 120
1       low    limit 160
router2#show queueing
Current fair queue configuration:
Current priority queue configuration:

List   Queue  Args
1       high   protocol ip          gt 1500
1       low    protocol ip          lt 256
1       normal protocol ip
1       normal interface Serial0
1       medium interface Ethernet0
1       high   limit 40
1       medium limit 80
1       normal limit 120
1       low    limit 160
Current custom queue configuration:
Current RED queue configuration:

List   Queue  Args
1       high   protocol ip          gt 1500
1       low    protocol ip          lt 256
1       normal protocol ip
1       normal interface Serial0
1       medium interface Ethernet0
1       high   limit 40

List   Queue  Args
1       medium limit 80
1       normal limit 120
1       low    limit 160
router2#
router2#


router2#show interface serial 0
Serial0 is up, line protocol is up
  Hardware is QUICC Serial
  Internet address is 172.16.3.2/24
  MTU 1500 bytes, BW 1544 Kbit, DLY 20000 usec, rely 255/255, load 1/255
  Encapsulation HDLC, loopback not set, keepalive set (10 sec)
  Last input 00:00:00, output 00:00:00, output hang never
  Last clearing of "show interface" counters never
  Input queue: 0/75/0 (size/max/drops); Total output drops: 0
  Queueing strategy: priority-list 1
  Output queue: high 0/40/0, medium 0/80/0, normal 0/120/0, low 0/160/0
  5 minute input rate 0 bits/sec, 0 packets/sec
  5 minute output rate 0 bits/sec, 0 packets/sec
     5304 packets input, 330350 bytes, 0 no buffer
     Received 1888 broadcasts, 0 runts, 0 giants
     3 input errors, 0 CRC, 3 frame, 0 overrun, 0 ignored, 0 abort
     5536 packets output, 350203 bytes, 0 underruns
     0 output errors, 0 collisions, 24 interface resets
     0 output buffer failures, 0 output buffers swapped out
     34 carrier transitions
     DCD=up  DSR=up  DTR=up  RTS=up  CTS=up

router2#show run
Building configuration...
```

```
Current configuration:
!
version 11.2
no service udp-small-servers
no service tcp-small-servers
!
hostname router2
!
enable secret 5 $1$MibE$cv2KLw/1kl8rmHtS6MBwe/
enable password 12345
!
!
interface Loopback0
 ip address 172.16.4.1 255.255.255.0
!
interface Loopback1
 ip address 172.16.5.1 255.255.255.0
!
interface Ethernet0
 no ip address
 shutdown
!
interface Serial0
 ip address 172.16.3.2 255.255.255.0
 priority-group 1
!
router eigrp 20
 network 172.16.0.0
!
no ip classless

priority-list 1 protocol ip high gt 1500
priority-list 1 protocol ip low lt 256
priority-list 1 protocol ip normal
priority-list 1 interface Serial0 normal
priority-list 1 interface Ethernet0 medium
priority-list 1 queue-limit 40 80 120 160
!
line con 0
line vty 0 4
 password 12345
 login
!
end

router2#
```

EIGRP Configuration Commands

The amount of bandwidth is controlled by two interface subcommands:

```
bandwidth <nnn>
```

and one of the following:

```
ip bandwidth-percent eigrp <AS-number>  <ppp>
appletalk eigrp-bandwidth-percent <ppp>
ipx bandwidth-percent eigrp <AS-number> <ppp>
```

for IP, AppleTalk, and IPX EIGRP, respectively.

The **bandwidth-percent** command tells EIGRP what percentage of the configured bandwidth it may use. The default is 50%. Since the **bandwidth** command is also used to set the routing protocol metric, it may be set to a particular value for policy reasons. The **bandwidth-percent** command can have values greater than 100 if the bandwidth is configured artificially low due to such policy reasons.

For example, the configuration:

```
interface Serial 0
bandwidth 56
ip bandwidth-percent eigrp 109 75
```

allows IP-EIGRP 109 to use 42Kbps (75% of 56Kbps) on Serial 0.

The configuration:

```
interface Serial 1
bandwidth 128
ipx bandwidth-percent eigrp 210 200
```

allows IPX-EIGRP 210 to use 256Kbps (200% of 128Kbps) on Serial 1.

Note: This assumes that Serial 1 is actually operating at a speed of at least 256Kbps.

Configuration Problems

If the bandwidth is configured to be a small value relative to the actual link speed, the enhanced implementation may converge at a slower rate than the earlier implementation. If the value is small enough and there are enough routes in the system, convergence may be so slow that it triggers "Stuck In Active" detection, which may prevent the network from ever converging. This state is evidenced by repeated messages of the form:

```
%DUAL-3-SIA: Route XXX stuck-in-active state in IP-EIGRP YY.  Cleaning up
```

The workaround for this problem is to raise the value of the "active" timer for EIGRP by configuring:

```
router eigrp
timers active-time
```

The default value in the enhanced code is three minutes; in earlier releases, the default is one minute. Raising this value would need to happen throughout the network.

If the bandwidth is configured to be too high (greater than the actual available bandwidth), the loss of EIGRP packets may occur. The packets will be retransmitted, but this may degrade convergence. The convergence in this case will be no slower than the earlier implementation, however.

Configuration Guidelines

Thse recommendations are described in terms of configuring the interface "bandwidth" parameter (with EIGRP being able to use 50% of that bandwidth by default). If the interface bandwidth configuration cannot be changed because of routing policy considerations, or for any other reason, the **bandwidth-percent** command should be used to control the EIGRP bandwidth. On low-speed interfaces, raising the available bandwidth for EIGRP above the default of 50% is advisable in order to improve convergence.

LAN Interfaces (Ethernet, Token Ring, FDDI)

The bandwidth parameter on LAN interfaces is set by default to the actual media speed, so no configuration should be necessary unless the bandwidth is explicitly configured to a very low value.

Point-to-Point Serial Interfaces (HDLC, PPP)

The bandwidth parameter defaults to T1 speed (1.544 Mbps) on serial interfaces. It should be set to the actual link speed.

NBMA Interfaces (Frame Relay, X.25, ATM)

It is particularly critical to configure NBMA interfaces correctly, because otherwise many EIGRP packets may be lost in the switched network. There are three basic rules:

1. The traffic that EIGRP is allowed to send on a single virtual circuit (VC) cannot exceed the capacity of that virtual circuit.
2. The total EIGRP traffic for all virtual circuits cannot exceed the access line speed of the interface.
3. The bandwidth allowed for EIGRP on each virtual circuit must be the same in each direction.

There are three different scenarios for NBMA interfaces.

- Pure Multipoint Configuration (no subinterfaces)
- Pure Point-to-Point Configuration (each VC on a separate subinterface)
- Hybrid Configuration (point-to-point and multipoint subinterfaces)

Each is examined separately.

Pure Multipoint Configuration (no subinterfaces)

In this configuration EIGRP will divide the configured bandwidth evenly across each virtual circuit. You must ensure that this will not overload each virtual circuit. For example, if you have a T1 access line with four 56K VCs, you should configure the bandwidth to be 224Kbps (4 * 56K) in order to avoid dropping packets. If the total bandwidth of the virtual circuits equals or exceeds the access line speed, configure the bandwidth to equal the access line speed. Note that if the virtual circuits are of different capacities, the bandwidth must be set to take into account the lowest capacity virtual circuit.

For instance, if a T1 access line has three 256K VCs and one 56K VC, the bandwidth should be set to 224Kbps (4 * 56K). In such configurations, putting at least the slow virtual circuit onto a point-to-point subinterface is strongly recommended (so that the bandwidth can be raised on the others).

Pure Point-to-Point Configuration (each VC on a separate subinterface)

This configuration allows maximum control, since the bandwidth can be configured separately on each subinterface, and is the best configuration if the virtual circuits have different capacities. Each subinterface bandwidth should be configured to be no greater than the available bandwidth on the associated virtual circuit, and the total bandwidth for all subinterfaces cannot exceed the available access line bandwidth. If the interface is oversubscribed, the access line bandwidth must be divided across each of the subinterfaces. For instance, if a T1 access line (1544 Kbps) has ten virtual circuits with a capacity of 256Kbps, the bandwidth on each subinterface should be configured to be 154Kbps (1544/10).

Hybrid Configuration (point-to-point and multipoint subinterfaces)

Hybrid configurations should use combinations of the two individual strategies, while ensuring that the three basic rules are followed.

Examples

The following examples illustrate the relationship between topology and configuration. Only the configuration commands pertaining to EIGRP's bandwidth usage are given.

Oversubscribed Hub-and-Spoke Frame Relay Configuration (Subinterfaces)

A fairly common configuration in networks with light amounts of transaction traffic is a hub-and-spoke configuration on which the access line to the hub is oversubscribed (since there is not usually enough data traffic to cause this to be a problem). In this scenario, assume a 256Kbps access line to the hub, with 56Kbps access lines to each of ten spoke sites. IP EIGRP process 123 is configured.

Because there is a maximum of 256Kbps available, we cannot allow any individual PVC to handle more than 25Kbps (256/10). Since this data rate is fairly low, and we don't expect very much user data traffic, we can allow EIGRP to use up to 90% of the bandwidth.

The hub configuration would look like:

```
interface Serial 0
encapsulation frame-relay

interface Serial 0.1 point-to-point
bandwidth 25
ip bandwidth-percent eigrp 123 90

interface Serial 0.2 point-to-point
bandwidth 25
ip bandwidth-percent eigrp 123 90
```

. . .

Each spoke router must be configured to limit EIGRP traffic to the same rate as that of the hub, in order to satisfy the third rule above. The spoke configuration would look like:

```
interface Serial 0
encapsulation frame-relay

interface Serial 0.1 point-to-point
bandwidth 25
ip bandwidth-percent eigrp 123 90
```

Note that EIGRP will not use more than 22.5Kbps (90% of 25K) on this interface, even though its capacity is 56Kbps. This configuration will not affect user data capacity, which will still be able to use the entire 56Kbps.

Alternatively, if you want to set the interface bandwidth to reflect the PVC capacity, you can adjust the bandwidth percentage for EIGRP. In this example, the desired bandwidth for EIGRP is (256K/10)*.9 = 23.04K; the bandwidth percentage would be 23.04K/56K = .41 (41%). So the same effect would be had by configuring:

```
interface Serial 0.1 point-to-point
bandwidth 56
ip bandwidth-percent eigrp 123 41
```

Full-mesh Frame Relay Configuration with Differing Access Line Speeds

In this configuration there is a fully-meshed Frame Relay network of four routers running IPX EIGRP process 456, configured as a multipoint network. Three of the four routers (routers A through C) have 256Kbps access lines, but one (router D) has only a 56Kbps access line. In this scenario, the configuration must restrict EIGRP's bandwidth in order to not overload the connection to router D. The simplest approach is to set the bandwidth to 56Kbps on all four routers:

```
interface Serial 0
encapsulation frame-relay
bandwidth 56
```

EIGRP will divide the bandwidth evenly across the three PVCs. Note, however, that this is overly restrictive for the PVCs connecting routers A through C, since they have sufficient capacity to handle a lot more traffic. One way of handling this situation is to convert the network to use point-to-point subinterfaces for all PVCs, as in the example above. Another way, which would require less configuration, is to break up the network by putting routers A through C on a fully meshed multipoint subinterface and making all of router D's connections point-to-point subinterfaces instead. The configuration on routers A through C would look like:

```
interface Serial 0
encapsulation frame-relay

interface Serial 0.1 multipoint
bandwidth 238

interface Serial 0.2 point-to-point
bandwidth 18
description PVC to Router D
```

Router D's configuration would look like:

```
interface Serial 0
encapsulation frame-relay

interface Serial 0.1 point-to-point
bandwidth 18
description PVC to Router A

interface Serial 0.2 point-to-point
bandwidth 18
description PVC to Router B

interface Serial 0.2 point-to-point
bandwidth 18
description PVC to Router C
```

Note that the multipoint subinterface is configured to 238 Kbps (256-18) and the point-to-point subinterfaces are configured to 18 Kbps (56/3).

Once again an alternative configuration can be used if it is desired to leave the "bandwidth" setting at its "natural" value. For the point-to-point interface, the desired bandwidth is (56K/3)*.5 = 9.33K; the percentage is 9.33K/56K = .16 (16%). For the multipoint interface

the desired bandwidth is (256K-18K)*.5 = 119K, so the bandwidth percent would be (119K/256K) = .46 (46%). The resulting configuration would be:

```
interface Serial 0.1 multipoint
bandwidth 256
ipx bandwidth-percent eigrp 456 46

interface Serial 0.2 point-to-point
bandwidth 56
description PVC to Router D
ipx bandwidth-percent eigrp 456 16
```

Basic EIGRP router configuration

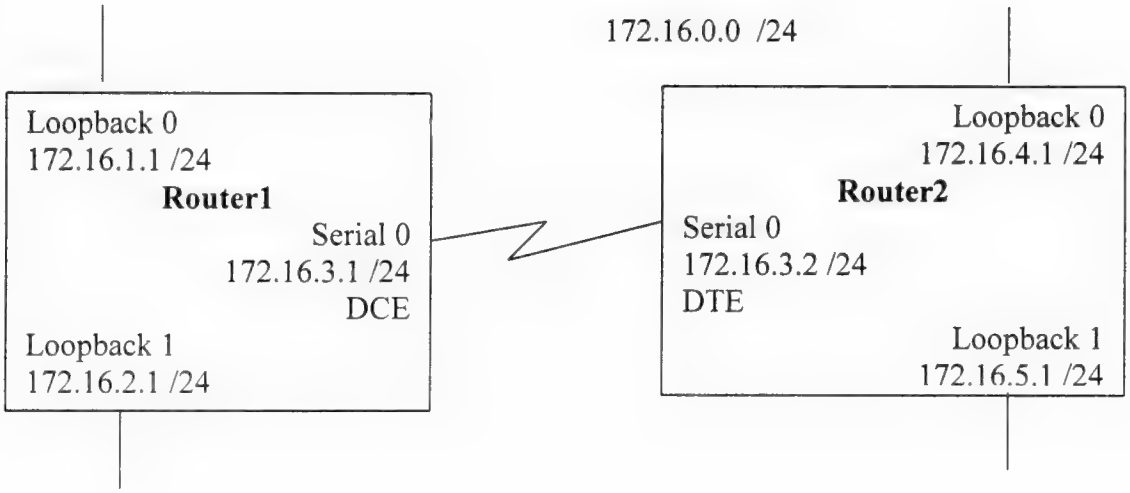

Config t
Hostname router1
Ip routing
Int loopback0
Ip address 172.16.1.1 255.255.255.0
Int loopback1
Ip address 172.16.2.1 255.255.255.0
Int serial0
Ip address 172.16.3.1 255.255.255.0
No shutdown
Clock rate 56000

Router eigrp 20
Network 172.16.1.0
Network 172.16.2.0
Network 172.16.3.0

Config t
hostname router2
ip routing
int loopback0
ip address 172.16.4.1 255.255.255.0
int loopback1
ip address 172.16.5.1 255.255.255.0
int serial0
ip address 172.16.3.2 255.255.255.0
no shutdown

router eigrp 20
network 172.16.4.0
network 172.16.5.0
network 172.16.3.0

```
router1#config t
Enter configuration commands, one per line.  End with CNTL/Z.
router1(config)#
router1(config)#router eigrp ?
  <1-65535>  Autonomous system number

router1(config)#router eigrp 20
router1(config-router)#network 172.16.1.0
router1(config-router)#network 172.16.2.0
router1(config-router)#network 172.16.3.0
router1(config-router)#
router1(config-router)#^Z
router1#
%SYS-5-CONFIG_I: Configured from console by console
router1#
router1#show running-config
Building configuration...

Current configuration:
!
hostname router1
!
enable secret 5 $1$8gsO$iFE02FDfh3ZbzZ0BujHPO/
enable password 12345
!
!
interface Loopback0
 ip address 172.16.1.1 255.255.255.0
!
interface Loopback1
 ip address 172.16.2.1 255.255.255.0
!
interface Ethernet0
 no ip address
 shutdown
!
!
interface Serial0
 ip address 172.16.3.1 255.255.255.0
 no fair-queue
 clockrate 56000
!
router eigrp 20
 network 172.16.0.0
!
line con 0
line aux 0
 transport input all
line vty 0 4
 password 12345
 login
!
end

router1#show ip route
Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, * - candidate default
```

```
Gateway of last resort is not set

     172.16.0.0 255.255.255.0 is subnetted, 5 subnets
D       172.16.4.0 [90/2297856] via 172.16.3.2, 00:07:32, Serial0
D       172.16.5.0 [90/2297856] via 172.16.3.2, 00:07:32, Serial0
C       172.16.1.0 is directly connected, Loopback0
C       172.16.2.0 is directly connected, Loopback1
C       172.16.3.0 is directly connected, Serial0
router1#
router1#

router1#config t
Enter configuration commands, one per line.  End with CNTL/Z.
router1(config)#
router1(config)#router eigrp 20
router1(config-router)#
router1(config-router)#  ?
Router configuration commands:
  auto-summary          Enable automatic network number summarization
  default-information   Control distribution of default information
  default-metric        Set metric of redistributed routes
  distance              Define an administrative distance
  distribute-list       Filter networks in routing updates
  eigrp                 EIGRP specific commands
  exit                  Exit from routing protocol configuration mode
  help                  Description of the interactive help system
  maximum-paths         Forward packets over multiple paths
  metric                Modify IGRP routing metrics and parameters
  neighbor              Specify a neighbor router
  network               Enable routing on an IP network
  no                    Negate or set default values of a command
  offset-list           Add or subtract offset from IGRP or RIP metrics
  passive-interface     Suppress routing updates on an interface
  redistribute          Redistribute information from another routing protocol
  timers                Adjust routing timers
  traffic-share         Algorithm for computing traffic share for alternate
                        routes
  variance              Control load balancing variance

router1(config-router)#

router1#show ip eigrp ?
  interfaces  IP-EIGRP interfaces
  neighbors   IP-EIGRP neighbors
  topology    IP-EIGRP Topology Table
  traffic     IP-EIGRP Traffic Statistics

router1#show ip eigrp interfaces
IP-EIGRP interfaces for process 20

                   Xmit Queue   Mean   Pacing Time   Multicast    Pending
Interface    Peers  Un/Reliable  SRTT   Un/Reliable   Flow Timer   Routes
Lo0            0        0/0        0        0/10          0           0
Lo1            0        0/0        0        0/10          0           0
Se0            1        0/0       36        0/15         50           0
```

```
router1#show ip eigrp interfaces detail
IP-EIGRP interfaces for process 20

                   Xmit Queue   Mean   Pacing Time   Multicast    Pending
Interface    Peers  Un/Reliable  SRTT   Un/Reliable   Flow Timer   Routes
Lo0            0        0/0        0       0/10          0           0
  Next xmit serial <none>
  Un/reliable mcasts: 0/0  Un/reliable ucasts: 0/0
  Mcast exceptions: 0  CR packets: 0  ACKs suppressed: 0
  Retransmissions sent: 0  Out-of-sequence rcvd: 0
Lo1            0        0/0        0       0/10          0           0
  Next xmit serial <none>
  Un/reliable mcasts: 0/0  Un/reliable ucasts: 0/0
  Mcast exceptions: 0  CR packets: 0  ACKs suppressed: 0
  Retransmissions sent: 0  Out-of-sequence rcvd: 0
Se0            1        0/0       36       0/15         50           0
  Next xmit serial <none>
  Un/reliable mcasts: 0/0  Un/reliable ucasts: 2/3
  Mcast exceptions: 0  CR packets: 0  ACKs suppressed: 0
  Retransmissions sent: 1  Out-of-sequence rcvd: 0


router1#
router1#show ip eigrp neighbors
IP-EIGRP neighbors for process 20
H   Address                 Interface   Hold Uptime   SRTT   RTO  Q  Seq
                                        (sec)         (ms)       Cnt Num
0   172.16.3.2              Se0           14 0:18:00    36   216  0  2
router1#
router1#show ip eigrp neighbors detail
IP-EIGRP neighbors for process 20
H   Address                 Interface   Hold Uptime   SRTT   RTO  Q  Seq
                                        (sec)         (ms)       Cnt Num
0   172.16.3.2              Se0           11 0:18:07    36   216  0  2
   Version 11.2/1.0, Retrans: 1, Retries: 0
router1#


router1#show ip eigrp topology
IP-EIGRP Topology Table for process 20

Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
       r - Reply status

P 172.16.4.0 255.255.255.0, 1 successors, FD is 2297856
         via 172.16.3.2 (2297856/128256), Serial0
P 172.16.5.0 255.255.255.0, 1 successors, FD is 2297856
         via 172.16.3.2 (2297856/128256), Serial0
P 172.16.1.0 255.255.255.0, 1 successors, FD is 128256
         via Connected, Loopback0
P 172.16.2.0 255.255.255.0, 1 successors, FD is 128256
         via Connected, Loopback1
P 172.16.3.0 255.255.255.0, 1 successors, FD is 2169856
         via Connected, Serial0
router1#
router1#show ip eigrp topology summary
IP-EIGRP Topology Table for process 20
Head serial 1, next serial 6
5 routes, 0 pending replies, 0 dummies
IP-EIGRP enabled on 3 interfaces, neighbors present on 1 interfaces
```

P metric

FD: feasible Distance
= metrics

```
Quiescent interfaces:  Se0
router1#
router1#show ip eigrp traffic
IP-EIGRP Traffic Statistics for process 20
  Hellos sent/received: 739/740
  Updates sent/received: 3/4
  Queries sent/received: 0/0
  Replies sent/received: 0/0
  Acks sent/received: 2/1
  Input queue high water mark 2, 0 drops
router1#
router1#show ip route 172.16.4.0
Routing entry for 172.16.4.0 255.255.255.0
  Known via "eigrp 20", distance 90, metric 2297856, type internal
  Redistributing via eigrp 20
  Last update from 172.16.3.2 on Serial0, 00:19:10 ago
  Routing Descriptor Blocks:
  * 172.16.3.2, from 172.16.3.2, 00:19:10 ago, via Serial0
      Route metric is 2297856, traffic share count is 1
      Total delay is 25000 microseconds, minimum bandwidth is 1544 Kbit
      Reliability 255/255, minimum MTU 1500 bytes
      Loading 1/255, Hops 1

router1#
router1#show ip protocol
Routing Protocol is "eigrp 20"
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set
  Default networks flagged in outgoing updates
  Default networks accepted from incoming updates
  EIGRP metric weight K1=1, K2=0, K3=1, K4=0, K5=0
  EIGRP maximum hopcount 100
  EIGRP maximum metric variance 1
  Redistributing: eigrp 20
  Automatic network summarization is in effect
  Routing for Networks:
    172.16.0.0
  Routing Information Sources:
    Gateway         Distance      Last Update
    172.16.3.2            90      0:19:23
  Distance: internal 90 external 170

router1#
router1#
```

VLSM (Variable Length Subnet Masking) Example using EIGRP

Using the previous EIGRP example we are going to change the length of the mask on the major interface From 172.16.0.0 /24 to use a variable length mask of /24 , /27,and /28 bit length.

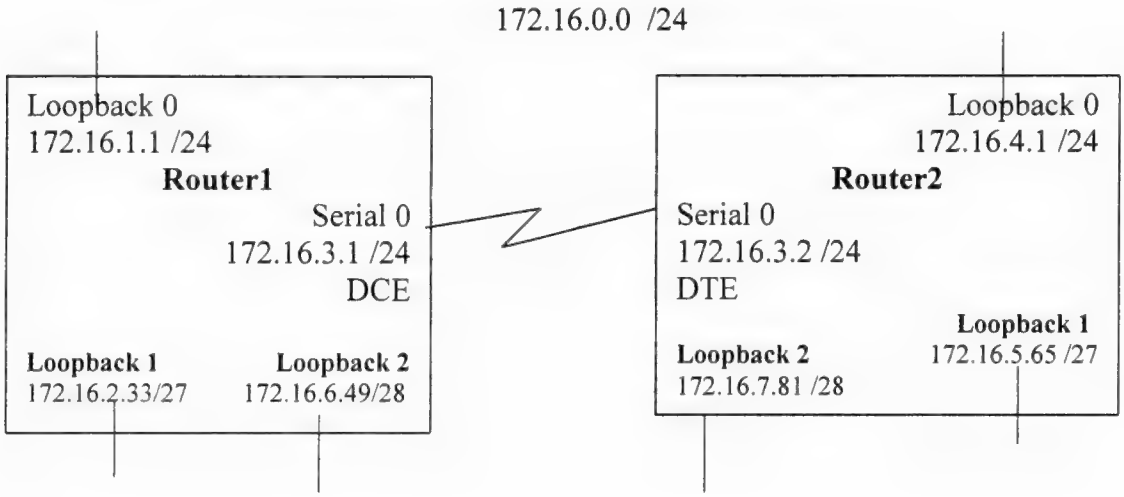

Router1 configuration:

1) *change loopback 1 interface ip address to 172.16.2.33 /27*
2) *add a loopback 2 interface with ip address 172.16.6.49 /28*
3) *start the router eigrp 20 process and add networks*
   172.16.2.32 and 172.16.6.48

Router2 configuration:

1) *change loopback 1 interface ip address to 172.16.5.65 /27*
2) *add a loopback 2 interface with ip address 172.16.7.81 /28*
3) *start the router eigrp 20 process and add networks*
   172.16.5.64 and 172.16.7.80

```
router1#config t
Enter configuration commands, one per line.  End with CNTL/Z.
router1(config)#interface loopback 1
router1(config-if)#ip address 172.16.2.33 255.255.255.224
router1(config-if)#
router1(config-if)#interface loopback 2
router1(config-if)#
%LINEPROTO-5-UPDOWN: Line protocol on Interface Loopback2, changed state to up
router1(config-if)#
%LINK-3-UPDOWN: Interface Loopback2, changed state to up
router1(config-if)#
router1(config-if)#ip address 172.16.6.49 255.255.255.240
router1(config-if)#
router1(config-if)#router eigrp 20
router1(config-router)#
router1(config-router)#network 172.16.2.32
router1(config-router)#network 172.16.6.48
router1(config-router)#^Z
router1#
```

```
%SYS-5-CONFIG_I: Configured from console by console
router1#

router2#
router2#config t

Enter configuration commands, one per line.  End with CNTL/Z.
router2(config)#interface loopback 1
router2(config-if)#ip address 172.16.5.65 255.255.255.224
router2(config-if)#
router2(config-if)#interface loopback 2
router2(config-if)#
%LINEPROTO-5-UPDOWN: Line protocol on Interface Loopback2, changed state to up
%LINK-3-UPDOWN: Interface Loopback2, changed state to up

router2(config-if)#ip address 172.16.7.81 255.255.255.240
router2(config-if)#
router2(config-if)#router eigrp 20
router2(config-router)#
router2(config-router)#network 172.16.5.64
router2(config-router)#network 172.16.7.80
router2(config-router)#
router2(config-router)#^Z
router2#
%SYS-5-CONFIG_I: Configured from console by console
router2#

router1#show running-config
Building configuration...

Current configuration:
!
version 11.0
service udp-small-servers
service tcp-small-servers
!
hostname router1
!
enable secret 5 $1$8gsO$iFE02FDfh3ZbzZ0BujHPO/
enable password 12345
!
!
interface Loopback0
 ip address 172.16.1.1 255.255.255.0
!
interface Loopback1
 ip address 172.16.2.33 255.255.255.224
!
interface Loopback2
 ip address 172.16.6.49 255.255.255.240
!
interface Ethernet0
 no ip address
 shutdown
!
interface Ethernet1
 no ip address
 shutdown
!
interface Serial0
```

```
 ip address 172.16.3.1 255.255.255.0
 no fair-queue
 clockrate 56000
!
interface Serial1
 no ip address
 shutdown
!
router eigrp 20
 network 172.16.0.0
!
!
line con 0
line aux 0
 transport input all
line vty 0 4
 password 12345
 login
!
end

router2#
router2#show run
Building configuration...

Current configuration:
!
version 11.2
no service udp-small-servers
no service tcp-small-servers
!
hostname router2
!
enable secret 5 $1$MibE$cv2KLw/1kl8rmHtS6MBwe/
enable password 12345
!
!
interface Loopback0
 ip address 172.16.4.1 255.255.255.0
!
interface Loopback1
 ip address 172.16.5.65 255.255.255.224
!
interface Loopback2
 ip address 172.16.7.81 255.255.255.240
!
interface Ethernet0
 no ip address
 shutdown
!
interface Serial0
 ip address 172.16.3.2 255.255.255.0
!
router eigrp 20
 network 172.16.0.0
!
no ip classless
!
line con 0
line vty 0 4
```

```
 password 12345
 login
!
end

router2#
router2#

router1#

router1#show ip route

Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, * - candidate default

Gateway of last resort is not set

     172.16.0.0 is variably subnetted, 7 subnets, 3 masks
C       172.16.6.48 255.255.255.240 is directly connected, Loopback2
C       172.16.2.32 255.255.255.224 is directly connected, Loopback1
D       172.16.4.0 255.255.255.0
           [90/2297856] via 172.16.3.2, 00:59:08, Serial0  [Admin Distance/metrica]
C       172.16.1.0 255.255.255.0 is directly connected, Loopback0
C       172.16.3.0 255.255.255.0 is directly connected, Serial0
D       172.16.7.80 255.255.255.240
           [90/2297856] via 172.16.3.2, 00:06:03, Serial0
D       172.16.5.64 255.255.255.224
           [90/2297856] via 172.16.3.2, 00:07:00, Serial0
router1#
router1#ping 172.16.5.65
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.5.65, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 36/37/44 ms
router1#
router1#show ip route 172.16.5.64
Routing entry for 172.16.5.64 255.255.255.224
  Known via "eigrp 20", distance 90, metric 2297856, type internal
  Redistributing via eigrp 20
  Last update from 172.16.3.2 on Serial0, 00:07:32 ago
  Routing Descriptor Blocks:
  * 172.16.3.2, from 172.16.3.2, 00:07:32 ago, via Serial0
      Route metric is 2297856, traffic share count is 1
      Total delay is 25000 microseconds, minimum bandwidth is 1544 Kbit
      Reliability 255/255, minimum MTU 1500 bytes
      Loading 1/255, Hops 1

router1#

router1#show ip eigrp topology

IP-EIGRP Topology Table for process 20

Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
       r - Reply status

P 172.16.6.48 255.255.255.240, 1 successors, FD is 128256
        via Connected, Loopback2
```

```
P 172.16.2.32 255.255.255.224, 1 successors, FD is 128256
         via Connected, Loopback1
P 172.16.4.0 255.255.255.0, 1 successors, FD is 2297856
         via 172.16.3.2 (2297856/128256), Serial0
P 172.16.1.0 255.255.255.0, 1 successors, FD is 128256
         via Connected, Loopback0
P 172.16.3.0 255.255.255.0, 1 successors, FD is 2169856
         via Connected, Serial0
P 172.16.7.80 255.255.255.240, 1 successors, FD is 2297856
         via 172.16.3.2 (2297856/128256), Serial0
P 172.16.5.64 255.255.255.224, 1 successors, FD is 2297856
         via 172.16.3.2 (2297856/128256), Serial0
router1#
router1#

router2#

router2#show ip route

Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, * - candidate default
       U - per-user static route, o - ODR

Gateway of last resort is not set

     172.16.0.0/16 is variably subnetted, 7 subnets, 3 masks
D       172.16.6.48/28 [90/2297856] via 172.16.3.1, 00:09:59, Serial0
D       172.16.2.32/27 [90/2297856] via 172.16.3.1, 00:10:51, Serial0
C       172.16.4.0/24 is directly connected, Loopback0
D       172.16.1.0/24 [90/2297856] via 172.16.3.1, 01:00:41, Serial0
C       172.16.3.0/24 is directly connected, Serial0
C       172.16.7.80/28 is directly connected, Loopback2
C       172.16.5.64/27 is directly connected, Loopback1
router2#
router2#ping 172.16.6.49

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.6.49, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 36/37/40 ms
router2#
router2#show ip route 172.16.6.48
Routing entry for 172.16.6.48/28
  Known via "eigrp 20", distance 90, metric 2297856, type internal
  Redistributing via eigrp 20
  Last update from 172.16.3.1 on Serial0, 00:10:38 ago
  Routing Descriptor Blocks:
  * 172.16.3.1, from 172.16.3.1, 00:10:38 ago, via Serial0
      Route metric is 2297856, traffic share count is 1
      Total delay is 25000 microseconds, minimum bandwidth is 1544 Kbit
      Reliability 255/255, minimum MTU 1500 bytes
      Loading 1/255, Hops 1
router2#
router2#show ip eigrp topology

IP-EIGRP Topology Table for process 20
Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
```

```
       r - Reply status
P 172.16.6.48/28, 1 successors, FD is 2297856
       via 172.16.3.1 (2297856/128256), Serial0
P 172.16.2.32/27, 1 successors, FD is 2297856
         via 172.16.3.1 (2297856/128256), Serial0
P 172.16.4.0/24, 1 successors, FD is 128256
         via Connected, Loopback0
P 172.16.1.0/24, 1 successors, FD is 2297856
         via 172.16.3.1 (2297856/128256), Serial0
P 172.16.3.0/24, 1 successors, FD is 2169856
         via Connected, Serial0
P 172.16.7.80/28, 1 successors, FD is 128256
         via Connected, Loopback2
P 172.16.5.64/27, 1 successors, FD is 128256
         via Connected, Loopback1
router2#
router2#
```

Redistributing Routing Protocols

While running a single routing protocol throughout your entire IP internetwork is desirable, multi-protocol routing is common for a number of reasons, including company mergers and multiple departments managed by multiple network administrators. Often, multi-protocols using redistribution is part of a network design. When you redistribute one protocol into another, remember that the metrics of each protocol play an important role in redistribution. Each protocol uses different metrics. For example, the RIP metric is based on hop count while the IGRP/EIGRP metric is based on bandwidth and delay. When routes are redistributed, you must define a metric that is understandable to the receiving protocol. There are two ways of defining metrics when redistributing routes.

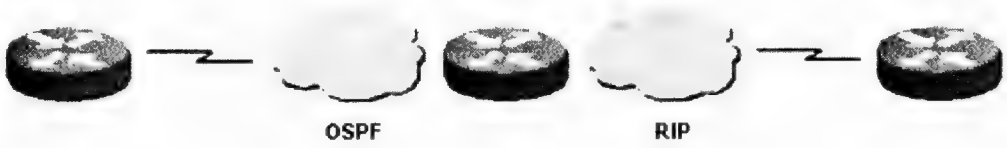

You can define the metric for that specific redistribution only:

```
router rip
redistribute static metric 1
redistribute ospf 1 metric 1
```

Or you can use the same metric as a default for all redistribution (Using the **default-metric** command saves work since it eliminates the need for defining the metric separately for each redistribution.):

```
router rip
redistribute static
redistribute ospf 1
default-metric 1
```

The following examples illustrate the syntax of redistribution. The examples show many protocols being redistributed between each other. We don't recommend running numerous protocols on one router, since it increases the risk of creating a routing loop or simply exhausting the router resources and CPU. However, there are some situations where a network administrator has to route more than one protocol and redistribute between them. The examples below demonstrate the syntax for handling this situation. Redistribution techniques for avoiding routing loops are discussed at the end of this document.

IGRP and EIGRP

The following output shows an IGRP/EIGRP router redistributing OSPF, RIP, and ISIS routes.

```
router igrp/eigrp 1
network 131.108.0.0
redistribute ospf 1
redistribute rip
redistribute isis
default-metric 10000 100 255 1 1500
```

IGRP and EIGRP need five metrics when redistributing other protocols: bandwidth, delay, reliability, load and MTU respectively. An example of IGRP metrics follows:

Metric	Value
bandwidth	10000 for Ethernet. Minimum bandwidth of the route may be 9.6.
delay	100 x 10 microseconds = 1 ms.
reliability	255 for 100% reliability.
loading	Effective bandwidth of the route in Kilobits per second (Kbps).
MTU	Minimum MTU of the router, usually equals the Ethernet bandwidth.

Multiple IGRP and EIGRP processes can run on the same router, with redistribution between them. For example, IGRP1 and IGRP2 can run on the same router. However, running two processes of the same protocol on the same router is rarely necessary, and it can consume the router's memory and CPU.

The redistribution of IGRP/EIGRP into another IGRP/EIGRP process doesn't require any metric conversion, so there is no need to define metrics or use the **default-metric** command during redistribution.

OSPF

The following output shows an OSPF router redistributing RIP, IGRP, EIGRP, and ISIS routes.

```
router ospf 1
network 131.108.0.0 0.0.255.255 area 0
redistribute rip metric 200 subnets
redistribute igrp 1 metric 100 subnets
redistribute eigrp 1 metric 100 subnets
redistribute isis metric 10 subnets
```

The OSPF metric is based on 10^8/ bandwidth of the link. For example, the OSPF cost of Ethernet is 10: $10^8/10^7 = 10$

Whenever there is a major net that is subnetted, you need to use the keyword *subnet* to redistribute protocols into OSPF. Without this keyword, OSPF only redistributes major nets that aren't subnetted. For instance, in the figure above, 131.108.0.0/16 is redistributed without the keyword *subnet*, but 141.108.100.0/24 isn't redistributed until we add the *subnet* keyword.

It's possible to run more than one OSPF process on the same router, but as mentioned before, running more than one process of the same protocol is rarely needed, and it consumes the router's memory and CPU.

You don't need to define metric or use the **default-metric** command when redistributing one OSPF process into another.

RIP

The following output shows a RIP router redistributing IGRP, EIGRP, OSPF, and ISIS routes.

```
router rip
network 131.108.0.0
redistribute igrp 1
redistribute eigrp 1
redistribute ospf 1
redistribute isis
default-metric 1
```

The RIP metric is composed of hop count, and the maximum valid metric is 15. Anything above 15 is considered infinite; you can use 16 to describe an infinite metric in RIP. When redistributing a protocol into RIP, we recommend using a low metric, such as 1. A high metric, such as 10, limits RIP even further: If we define a metric of 10 for redistributed routes, these routes can only be advertised to routers up to 5 hops away, at which point the metric (hopcount) exceeds 15. By defining a metric of 1, you enable a route to travel the maximum number of hops in a RIP domain.

ISIS

The following output shows an ISIS router redistributing RIP, IGRP, EIGRP, and OSPF routes.

```
router isis
net ..00
redistribute rip metric 20
redistribute igrp 1 metric 20
redistribute eigrp 1 metric 20
redistribute ospf 1 metric 20
```

The ISIS metric needs to be between 1 and 63. There is no default-metric option in ISIS, so you should define a metric for each protocol, as shown in the example above.

How to Avoid Oscillation Due To Feedback in IGRP/EIGRP

On a given network, you should only redistribute external routes in the routing process. The following example allows only non-131.108.0.0 network information to be redistributed into an IGRP 1 process.

```
router igrp 1
redistribute rip metric 1000 20000 255 1 12000
network 131.108.0.0
distribute-list 1 out rip

access-list 1 deny 131.108.0.0 0.0.255.255
access-list 1 permit 0.0.0.0 255.255.255.255
```

Similarly, you can filter updates sent from a routing process or from a process on a given interface, as follows.

```
router rip
redistribute igrp 109 metric 1
network 131.108.0.0
distribute-list 1 out interface e0
```

Or, you can establish a passive interface to avoid sending routing updates on a given interface.

```
distribute-list 2 out igrp 109
passive-interface ethernet 0

access-list 2 permit 131.108.0.0
access-list 2 deny 0.0.0.0 255.255.255.255
```

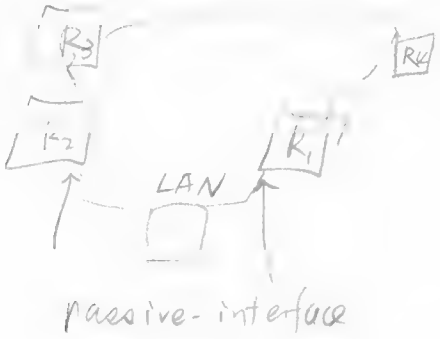

Controlling routing updates and traffic management

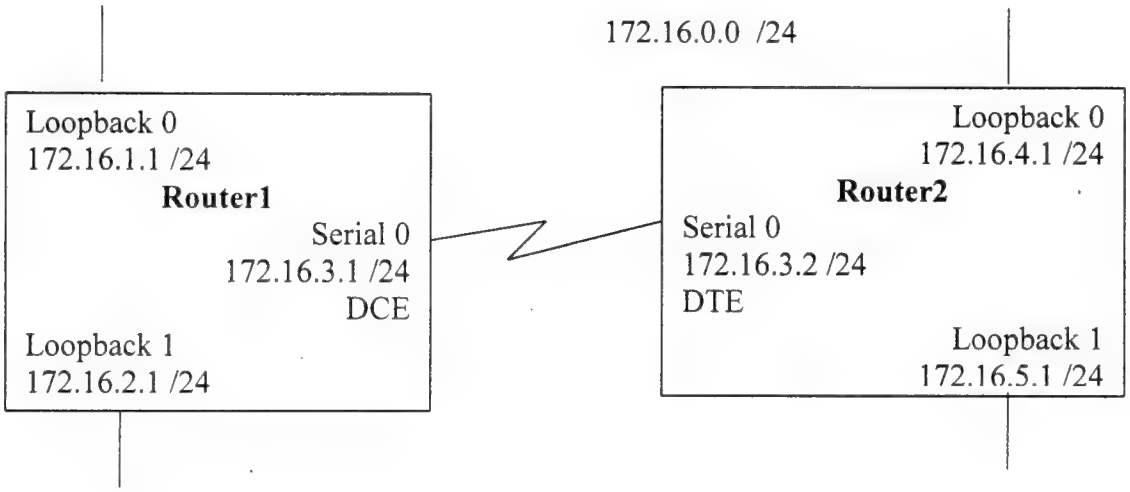


```
Config t
Hostname router1
Ip routing
Int loopback0
Ip address 172.16.1.1 255.255.255.0
Int loopback1
Ip address 172.16.2.1 255.255.255.0
Int serial0
Ip address 172.16.3.1 255.255.255.0
No shutdown
Clock rate 56000

Router eigrp 20
Network 172.16.1.0
Network 172.16.2.0
Network 172.16.3.0
```

```
Config t
hostname router2
ip routing
int loopback0
ip address 172.16.4.1 255.255.255.0
int loopback1
ip address 172.16.5.1 255.255.255.0
int serial0
ip address 172.16.3.2 255.255.255.0
no shutdown

router eigrp 20
network 172.16.4.0
network 172.16.5.0
network 172.16.3.0
```

Using a previous EIGRP routing configuration, we want to control routing updates from router1 to router2 (and vice-versus) thru the serial link. There are two methods which we will explore.

1) passive-interface
2) access-lists and distribute-list

The first method using passive-interface will block ***all*** *routing updates from router1 to router2.*

Notice that router2 knows about all networks attached to router1 before we start.

```
router2#show ip route
Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, * - candidate default
       U - per-user static route, o - ODR

Gateway of last resort is not set

     172.16.0.0/24 is subnetted, 5 subnets
C       172.16.4.0 is directly connected, Loopback0
C       172.16.5.0 is directly connected, Loopback1
D       172.16.1.0 [90/2297856] via 172.16.3.1, 00:02:53, Serial0
D       172.16.2.0 [90/2297856] via 172.16.3.1, 00:02:53, Serial0
C       172.16.3.0 is directly connected, Serial0
router2#
```

Next we will place router1's serial interface in the passive mode so that it will no longer send out route updates of any kind.

```
router1#
router1#config t
Enter configuration commands, one per line.  End with CNTL/Z.
router1(config)#router eigrp 20
router1(config-router)#
router1(config-router)# ?
Router configuration commands:
  auto-summary          Enable automatic network number summarization
  default-information   Control distribution of default information
  default-metric        Set metric of redistributed routes
  distance              Define an administrative distance
  distribute-list       Filter networks in routing updates
  eigrp                 EIGRP specific commands
  exit                  Exit from routing protocol configuration mode
  help                  Description of the interactive help system
  maximum-paths         Forward packets over multiple paths
  metric                Modify IGRP routing metrics and parameters
  neighbor              Specify a neighbor router
  network               Enable routing on an IP network
  no                    Negate or set default values of a command
  offset-list           Add or subtract offset from IGRP or RIP metrics
  passive-interface     Suppress routing updates on an interface
  redistribute          Redistribute information from another routing protocol
```

```
  timers               Adjust routing timers
  traffic-share        Algorithm for computing traffic share for alternate
                       routes
  variance             Control load balancing variance

router1(config-router)#passive-interface serial 0
router1(config-router)#
router1(config-router)#^Z
router1#
%SYS-5-CONFIG_I: Configured from console by console
router1#
```

Notice that router2 no longer knows about router1's attached networks.

```
router2#show ip route
Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, * - candidate default
       U - per-user static route, o - ODR

Gateway of last resort is not set

     172.16.0.0/24 is subnetted, 3 subnets
C       172.16.4.0 is directly connected, Loopback0
C       172.16.5.0 is directly connected, Loopback1
C       172.16.3.0 is directly connected, Serial0
router2#
```

we can no longer ping router1 !!!

```
router2#ping 172.16.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.1.1, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
router2#
```

*One way to fix this is to configure a **static default route** on router2 pointing to router1.*

```
router2#config t
Enter configuration commands, one per line.  End with CNTL/Z.
router2(config)#ip ?
Global IP configuration subcommands:
  accounting-list       Select hosts for which IP accounting information is
                        kept
  accounting-threshold  Sets the maximum number of accounting entries
  accounting-transits   Sets the maximum number of transit entries
  address-pool          Specify default IP address pooling mechanism
  alias                 Alias an IP address to a TCP port
  bootp                 Config BOOTP services
  classless             Follow classless routing forwarding rules
```

static Default Route

① ip route 0.0.0.0 0.0.0.0 S0/own interface

② ip route 0.0.0.0 0.0.0.0 172.36.1.2
next hop

③ ip default-network S0
own interface

④ ip default-network 172.16.3.2
ip address of next hop

```
  default-gateway       Specify default gateway (if not routing IP)
  default-network       Flags networks as candidates for default routes
  dhcp-server           Specify address of DHCP server to use
  domain-list           Domain name to complete unqualified host names.
  domain-lookup         Enable IP Domain Name System hostname translation
  domain-name           Define the default domain name
  dvmrp                 DVMRP global commands
  flow-export           Specify host/port to send flow statistics
  forward-protocol      Controls forwarding of physical and directed IP
                        broadcasts
  gratuitous-arps       Generate gratuitous ARPs for PPP/SLIP peer addresses
  host                  Add an entry to the ip hostname table
  host-routing          Enable host-based routing (proxy ARP and redirect)

router2(config)#ip default-network 172.16.3.0
router2(config)#^Z
router2#
%SYS-5-CONFIG_I: Configured from console by console
router2#ping 172.16.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.1.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 36/36/40 ms
router2#

router2#show ip route
Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, * - candidate default
       U - per-user static route, o - ODR

Gateway of last resort is not set

     172.16.0.0/16 is variably subnetted, 4 subnets, 2 masks
C       172.16.4.0/24 is directly connected, Loopback0
C       172.16.5.0/24 is directly connected, Loopback1
S       172.16.0.0/16 [1/0] via 172.16.3.0
C       172.16.3.0/24 is directly connected, Serial0
router2#
```

Now remove the static default-network route and remove the passive-interface statement before proceeding to the next method !!!!

Router2

Config t
No ip route 172.16.0.0 255.255.0.0 172.16.3.0

Router1
Config t
Router eigrp 20
No passive-interface serial 0

Verify that both routers know about all networks !!!

*The passive-interface method blocked **all** routing updates. The **distribute-list method can block only certain route updates**. We want to block only the route update for network 172.16.1.0 from router1 to router2.*

```
router1#config t
Enter configuration commands, one per line.  End with CNTL/Z.
router1(config)#
router1(config)#access-list 10 deny 172.16.1.0 0.0.0.255
router1(config)#access-list 10 permit any
router1(config)#
router1(config)#router eigrp 20
router1(config-router)# ?
Router configuration commands:
  auto-summary          Enable automatic network number summarization
  default-information   Control distribution of default information
  default-metric        Set metric of redistributed routes
  distance              Define an administrative distance
  distribute-list       Filter networks in routing updates
  eigrp                 EIGRP specific commands
  exit                  Exit from routing protocol configuration mode
  help                  Description of the interactive help system
  maximum-paths         Forward packets over multiple paths
  metric                Modify IGRP routing metrics and parameters
  neighbor              Specify a neighbor router
  network               Enable routing on an IP network
  no                    Negate or set default values of a command
  offset-list           Add or subtract offset from IGRP or RIP metrics
  passive-interface     Suppress routing updates on an interface
  redistribute          Redistribute information from another routing protocol
  timers                Adjust routing timers
  traffic-share         Algorithm for computing traffic share for alternate
                        routes
  variance              Control load balancing variance

router1(config-router)#distribute-list ?
  <1-199>  A standard IP access list number

router1(config-router)#distribute-list 10 ?
  in   Filter incoming routing updates
  out  Filter outgoing routing updates

router1(config-router)#distribute-list 10 out ?
  Ethernet    IEEE 802.3
  Loopback    Loopback interface
  Null        Null interface
  Serial      Serial
  bgp         Border Gateway Protocol (BGP)
  connected   Connected
  egp         Exterior Gateway Protocol (EGP)
  eigrp       Enhanced Interior Gateway Routing Protocol (EIGRP)
  igrp        Interior Gateway Routing Protocol (IGRP)
  isis        ISO IS-IS
  iso-igrp    IGRP for OSI networks
  ospf        Open Shortest Path First (OSPF)
  rip         Routing Information Protocol (RIP)
  static      Static routes
```

```
  <cr>

router1(config-router)#distribute-list 10 out serial 0
router1(config-router)#
router1(config-router)#^Z
router1#
%SYS-5-CONFIG_I: Configured from console by console
router1#

router1#show run
Building configuration...

Current configuration:
!
version 11.0
service udp-small-servers
service tcp-small-servers
!
hostname router1
!
enable secret 5 $1$8gsO$iFE02FDfh3ZbzZ0BujHPO/
enable password 12345
!
!
interface Loopback0
 ip address 172.16.1.1 255.255.255.0
!
interface Loopback1
 ip address 172.16.2.1 255.255.255.0
!
interface Ethernet0
 no ip address
 shutdown
!
interface Ethernet1
 no ip address
 shutdown
!
interface Serial0
 ip address 172.16.3.1 255.255.255.0
 no fair-queue
 clockrate 56000
!
interface Serial1
 no ip address
 shutdown
!
router eigrp 20
 network 172.16.0.0
 distribute-list 10 out Serial0
!
access-list 10 deny   172.16.1.0 0.0.0.255
access-list 10 permit any
!
line con 0
line aux 0
 transport input all
line vty 0 4
 password 12345
 login
!
end

router1#show access-list
Standard IP access list 10
    deny   172.16.1.0, wildcard bits 0.0.0.255
    permit any
```

```
router1#
router1#show ip protocol
Routing Protocol is "eigrp 20"
  Outgoing update filter list for all interfaces is not set
    Serial0 filtered by 10
  Incoming update filter list for all interfaces is not set
  Default networks flagged in outgoing updates
  Default networks accepted from incoming updates
  EIGRP metric weight K1=1, K2=0, K3=1, K4=0, K5=0
  EIGRP maximum hopcount 100
  EIGRP maximum metric variance 1
  Redistributing: eigrp 20
  Automatic network summarization is in effect
  Routing for Networks:
    172.16.0.0
  Routing Information Sources:
    Gateway         Distance      Last Update
    172.16.3.2            90      0:00:21
  Distance: internal 90 external 170

router1#
```

Notice that on router2 that the route to network 172.16.1.0 in no longer visible but the route to 172.16.2.0 is still available.

```
router2#
router2# show ip route
Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, * - candidate default
       U - per-user static route, o - ODR

Gateway of last resort is not set

     172.16.0.0/24 is subnetted, 4 subnets
C       172.16.4.0 is directly connected, Loopback0
C       172.16.5.0 is directly connected, Loopback1
D       172.16.2.0 [90/2297856] via 172.16.3.1, 00:00:50, Serial0
C       172.16.3.0 is directly connected, Serial0
router2#
```

We can no longer ping 172.16.1.1 but can still access 172.16.2.1 !!!

```
router2#ping 172.16.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.1.1, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
router2#
router2#ping 172.16.2.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.2.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 36/36/36 ms
router2#
```

Route Redistribution Exercise RIP to IGRP

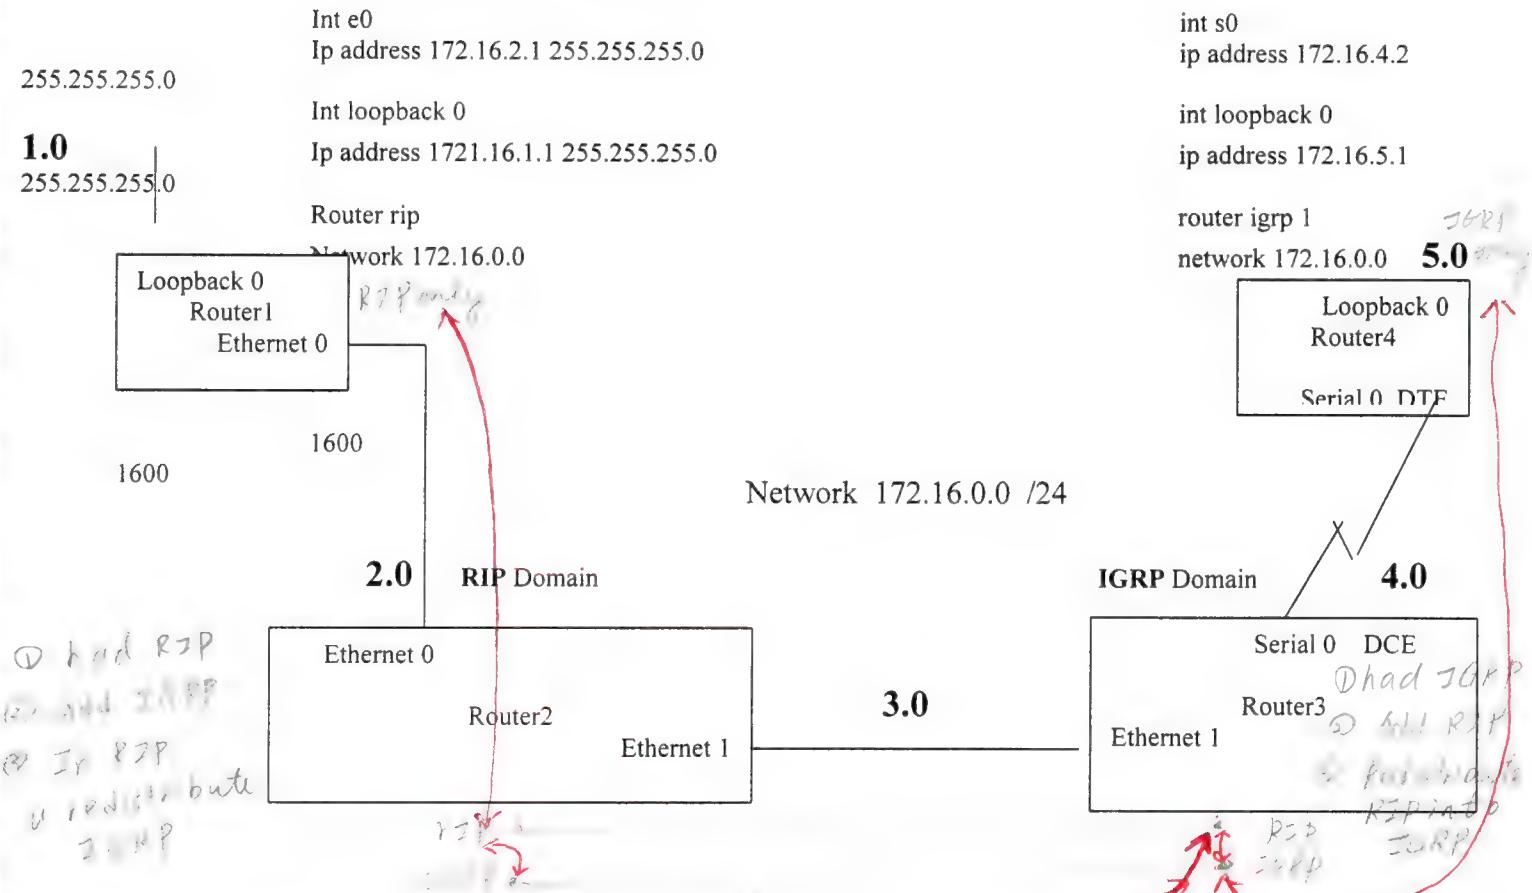

int e0
ip address 172.16.2.2 255.255.255.0
255.255.255.0
int e1
ip address 172.16.3.1 255.255.255.0
255.255.255.0
router rip
network 172.16.0.0

int e1
ip address 172.16.3.2

int s0
ip address 172.16.4.1

clock rate 56000
router igrp 1
network 172.16.0.0

```
router1#sh run
Building configuration...

Current configuration:
!
version 11.2
no service udp-small-servers
no service tcp-small-servers
!
hostname router1
!
enable secret 5 $1$gcTi$f90oMOedGU9Q13YwPOf6y.
enable password 12345
!
!
interface Loopback0
 ip address 172.16.1.1 255.255.255.0
!
interface Ethernet0
 ip address 172.16.2.1 255.255.255.0
!
interface Serial0
 no ip address
 shutdown
!
router rip
 network 172.16.0.0
!
no ip classless
!
line con 0
line vty 0 4
 password 12345
 login
!
end

router1#
router1#

router2#
router2#sh run
Building configuration...

Current configuration:
!
version 11.2
!
hostname router2
!
enable secret 5 $1$F2sX$hKY48JDuhyW2cmiW.1Huw.
enable password 12345
!
!
interface Ethernet0
 ip address 172.16.2.2 255.255.255.0
!
interface Ethernet1
```

```
 ip address 172.16.3.1 255.255.255.0
!
interface Serial0
 no ip address
 shutdown
!
interface Serial1
 no ip address
 shutdown
!
router rip
 network 172.16.0.0
!
no ip classless
!
line con 0
line aux 0
line vty 0 4
 password 12345
 login
!
end

router2#
router2#

router3#
router3#sh run
Building configuration...

Current configuration:
!
version 11.0
service udp-small-servers
service tcp-small-servers
!
hostname router3
!
enable secret 5 $1$Np2t$c83TCa2O8dKG9Wl4MshaS1
enable password 12345
!
!
interface Ethernet0
 no ip address
 shutdown
!
interface Ethernet1
 ip address 172.16.3.2 255.255.255.0
!
interface Serial0
 ip address 172.16.4.1 255.255.255.0
 no fair-queue
 clockrate 56000
!
interface Serial1
 no ip address
 shutdown
!
router igrp 1
 network 172.16.0.0
!
```

```
!
line con 0
line aux 0
 transport input all
line vty 0 4
 password 12345
 login
!
end

router3#

router4#
router4#sh run
Building configuration...

Current configuration:
!
version 11.2
no service udp-small-servers
no service tcp-small-servers
!
hostname router4
!
enable secret 5 $1$Ixbz$.G9chMxpTZrjNpsRPYqa2/
enable password 12345
!
!
interface Loopback0
 ip address 172.16.5.1 255.255.255.0
!
interface Ethernet0
 no ip address
 shutdown
!
interface Serial0
 ip address 172.16.4.2 255.255.255.0
!
router igrp 1
 network 172.16.0.0
!
no ip classless
!
line con 0
line vty 0 4
 password 12345
 login
!
end

router4#
router4#
```

Looking at the route tables for each router you will see the following:
Router1 and Router2 do not know about subnets 4.0 and 5.0
Router3 and Router4 do not know about subnets 1.0 and 2.0
Because the two sets of routers are running different routing protocols we must perform route redistribution to inject the other routes into each routing domain

```
router1#show ip route
Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, * - candidate default
       U - per-user static route, o - ODR

Gateway of last resort is not set

     172.16.0.0/24 is subnetted, 3 subnets
C       172.16.1.0 is directly connected, Loopback0
C       172.16.2.0 is directly connected, Ethernet0
R       172.16.3.0 [120/1] via 172.16.2.2, 00:00:09, Ethernet0
router1#
router1#
router2#
router2#show ip route
Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, * - candidate default
       U - per-user static route, o - ODR

Gateway of last resort is not set

     172.16.0.0/24 is subnetted, 3 subnets
R       172.16.1.0 [120/1] via 172.16.2.1, 00:00:17, Ethernet0
C       172.16.2.0 is directly connected, Ethernet0
C       172.16.3.0 is directly connected, Ethernet1
router2#
router2#

router3#
router3#show ip route
Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, * - candidate default

Gateway of last resort is not set

     172.16.0.0 255.255.255.0 is subnetted, 3 subnets
C       172.16.4.0 is directly connected, Serial0
I       172.16.5.0 [100/8976] via 172.16.4.2, 00:00:41, Serial0
C       172.16.3.0 is directly connected, Ethernet1
router3#
router3#

router4#
router4#show ip route
Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
```

```
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, * - candidate default
       U - per-user static route, o - ODR

Gateway of last resort is not set

     172.16.0.0/24 is subnetted, 3 subnets
C       172.16.4.0 is directly connected, Serial0
C       172.16.5.0 is directly connected, Loopback0
I       172.16.3.0 [100/8576] via 172.16.4.1, 00:00:34, Serial0
router4#
router4#
router4#
```

To perform route redistribution we will enable IGRP on router2 and then redistribute these routes by injecting them into the RIP domain

Notice once again that router2 does not know about subnets 4.0 and 5.0

```
router2#show ip route
Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, * - candidate default
       U - per-user static route, o - ODR

Gateway of last resort is not set

     172.16.0.0/24 is subnetted, 3 subnets
R       172.16.1.0 [120/1] via 172.16.2.1, 00:00:11, Ethernet0
C       172.16.2.0 is directly connected, Ethernet0
C       172.16.3.0 is directly connected, Ethernet1
router2#
router2#config t
Enter configuration commands, one per line.  End with CNTL/Z.
router2(config)#router igrp 1
router2(config-router)#network 172.16.0.0
router2(config-router)#
router2(config-router)#^Z
router2#
```

Notice now that router2 has learned about all subnets via both routing processes because router2 is running both routing protocols now.

```
router2#show ip route
Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, * - candidate default
       U - per-user static route, o - ODR

Gateway of last resort is not set

     172.16.0.0/24 is subnetted, 5 subnets
I       172.16.4.0 [100/8576] via 172.16.3.2, 00:00:03, Ethernet1
I       172.16.5.0 [100/9076] via 172.16.3.2, 00:00:03, Ethernet1
R       172.16.1.0 [120/1] via 172.16.2.1, 00:00:26, Ethernet0
```

```
C       172.16.2.0 is directly connected, Ethernet0
C       172.16.3.0 is directly connected, Ethernet1
router2#
router2#
```

Notice that router1 still has not learned about all of the subnets yet.

```
router1#
router1#show ip route
Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, * - candidate default
       U - per-user static route, o - ODR

Gateway of last resort is not set

     172.16.0.0/24 is subnetted, 3 subnets
C       172.16.1.0 is directly connected, Loopback0
C       172.16.2.0 is directly connected, Ethernet0
R       172.16.3.0 [120/1] via 172.16.2.2, 00:00:00, Ethernet0
```

In order for router1 to learn about all of the possible subnets we need the router2 RIP process to redistribute the IGRP routes it has learned to router1.

```
router2#
router2#config t
Enter configuration commands, one per line.  End with CNTL/Z.
router2(config)#router rip
router2(config-router)#
router2(config-router)#?
Router configuration commands:
  auto-summary            Enable automatic network number summarization
  default                 Set a command to its defaults
  default-metric          Set metric of redistributed routes
  distance                Define an administrative distance
  distribute-list         Filter networks in routing updates
  exit                    Exit from routing protocol configuration mode
  help                    Description of the interactive help system
  maximum-paths           Forward packets over multiple paths
  neighbor                Specify a neighbor router
  network                 Enable routing on an IP network
  no                      Negate a command or set its defaults
  offset-list             Add or subtract offset from IGRP or RIP metrics
  output-delay            Interpacket delay for RIP updates
  passive-interface       Suppress routing updates on an interface
  redistribute            Redistribute information from another routing
                          protocol
  timers                  Adjust routing timers
  validate-update-source  Perform sanity checks against source address of
                          routing updates
  version                 Set routing protocol version
```

```
router2(config-router)#redistribute ?
  bgp         Border Gateway Protocol (BGP)
  connected   Connected
  egp         Exterior Gateway Protocol (EGP)
  eigrp       Enhanced Interior Gateway Routing Protocol (EIGRP)
  igrp        Interior Gateway Routing Protocol (IGRP)
  isis        ISO IS-IS
  iso-igrp    IGRP for OSI networks
  mobile      Mobile routes
  odr         On Demand stub Routes
  ospf        Open Shortest Path First (OSPF)
  rip         Routing Information Protocol (RIP)
  static      Static routes ( Administration . 0) multiple

router2(config-router)#redistribute igrp 1
router2(config-router)#
```

Next we need to reset the metrics so that RIP will understand the redistributed or injected IGRP routes.

```
router2(config-router)#default-metric ?
  <1-4294967295>  Default metric

router2(config-router)#default-metric 3   no very high metric
router2(config-router)#
router2(config-router)#^Z
router2#
```

*Now you will see that **router1 has learned about subnets 4.0 and 5.0** via RIP which were redistributed with a hop count metric of 3.*

```
router1#
router1#show ip route
Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, * - candidate default
       U - per-user static route, o - ODR

     172.16.0.0/24 is subnetted, 5 subnets
R       172.16.4.0 [120/3] via 172.16.2.2, 00:00:02, Ethernet0
R       172.16.5.0 [120/3] via 172.16.2.2, 00:00:02, Ethernet0
C       172.16.1.0 is directly connected, Loopback0
C       172.16.2.0 is directly connected, Ethernet0
R       172.16.3.0 [120/1] via 172.16.2.2, 00:00:02, Ethernet0
router1#
router1#ping 172.16.5.1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.5.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 36/39/44 ms
```

Router1 can successfully communicate with the far network !!!!

Router3 and Router4 still do not know about subnet 1.0 so the last step is to configure Router3 for route redistribution

```
router3#show ip route
Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, * - candidate default

Gateway of last resort is not set

     172.16.0.0 255.255.255.0 is subnetted, 4 subnets
C       172.16.4.0 is directly connected, Serial0
I       172.16.5.0 [100/8976] via 172.16.4.2, 00:00:48, Serial0
I       172.16.2.0 [100/1200] via 172.16.3.1, 00:00:07, Ethernet1
C       172.16.3.0 is directly connected, Ethernet1
```

The first step is start the RIP routing process on router3

```
router3#
router3#config t
Enter configuration commands, one per line.  End with CNTL/Z.
router3(config)#router rip
router3(config-router)#network 172.16.0.0
router3(config-router)#^Z
router3#
router3#show ip route
Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, * - candidate default

Gateway of last resort is not set

     172.16.0.0 255.255.255.0 is subnetted, 5 subnets
C       172.16.4.0 is directly connected, Serial0
I       172.16.5.0 [100/8976] via 172.16.4.2, 00:00:08, Serial0
R       172.16.1.0 [120/2] via 172.16.3.1, 00:00:15, Ethernet1
I       172.16.2.0 [100/1200] via 172.16.3.1, 00:00:55, Ethernet1
C       172.16.3.0 is directly connected, Ethernet1
router3#
```

Router3 has learned about all subnets but Router4 still has not !!

```
router4#
router4#show ip route
Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, * - candidate default
       U - per-user static route, o - ODR

Gateway of last resort is not set

     172.16.0.0/24 is subnetted, 4 subnets
C       172.16.4.0 is directly connected, Serial0
C       172.16.5.0 is directly connected, Loopback0
I       172.16.2.0 [100/8676] via 172.16.4.1, 00:01:08, Serial0
I       172.16.3.0 [100/8576] via 172.16.4.1, 00:01:08, Serial0
```

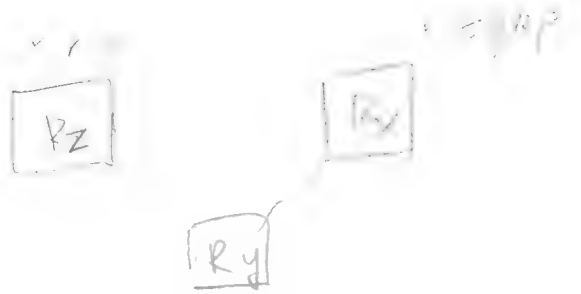

router RIP
net —
redistribute IGRP 10

router igrp 10
net —
redistribute rip
default-metric

The next step is to the IGRP routes to the RIP domain, but IGRP metrics are different than that of RIP. IGRP metrics use the following:
IGRP uses ***Bandwidth Delay Reliability Load MTU***
These parameters can be gotten from the ethernet interface

```
router3#
router3#show interface e1
Ethernet1 is up, line protocol is up
  Hardware is Lance, address is 0010.7b3b.f8c1 (bia 0010.7b3b.f8c1)
  Internet address is 172.16.3.2 255.255.255.0
  MTU 1500 bytes, BW 10000 Kbit, DLY 1000 usec, rely 255/255, load 1/255
  Encapsulation ARPA, loopback not set, keepalive set (10 sec)
  ARP type: ARPA, ARP Timeout 4:00:00
  Last input 0:00:04, output 0:00:01, output hang never
  Last clearing of "show interface" counters never
  Output queue 0/40, 0 drops; input queue 0/75, 0 drops
  5 minute input rate 0 bits/sec, 0 packets/sec
  5 minute output rate 0 bits/sec, 0 packets/sec
     317 packets input, 44654 bytes, 0 no buffer
     Received 304 broadcasts, 0 runts, 0 giants
     0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored, 0 abort
     0 input packets with dribble condition detected
     735 packets output, 72883 bytes, 0 underruns
     0 output errors, 0 collisions, 1 interface resets, 0 restarts
     0 output buffer failures, 0 output buffers swapped out
router3#
```

Now we can redistribute the routes.

```
router3#config t
Enter configuration commands, one per line.  End with CNTL/Z.
router3(config)#router igrp 1
router3(config-router)#redistribute rip
router3(config-router)#default-metric ?
  <1-4294967295>  IGRP bandwidth metric in kilobits per second

router3(config-router)#default-metric 10000 ?
  <0-4294967295>  IGRP delay metric, in 10 microsecond units

router3(config-router)#default-metric 10000 100 ?
  <0-255>  IGRP reliability metric where 255 is 100% reliable

router3(config-router)#default-metric 10000 100 255 ?
  <1-255>  IGRP Effective bandwidth metric (Loading) where 255 is 100% loaded

router3(config-router)#default-metric 10000 100 255 1 ?
  <1-4294967295>  IGRP Maximum Transmission Unit metric of the path

router3(config-router)#default-metric 10000 100 255 1 1500
router3(config-router)#
router3(config-router)#^Z
router3#
```

Now notice that router4 has learned about all routes and we can ping the far network interface !!!

```
router4#
router4#show ip route
Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
```

```
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, * - candidate default
       U - per-user static route, o - ODR


     172.16.0.0/24 is subnetted, 5 subnets
C       172.16.4.0 is directly connected, Serial0
C       172.16.5.0 is directly connected, Loopback0
I       172.16.1.0 [100/8576] via 172.16.4.1, 00:00:10, Serial0
I       172.16.2.0 [100/8676] via 172.16.4.1, 00:00:10, Serial0
I       172.16.3.0 [100/8576] via 172.16.4.1, 00:00:10, Serial0
router4#
router4#ping 172.16.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.1.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 36/39/44 ms
router4#
```

Configuring Weighted Fair Queuing and Priority Queuing

Weighted Fair Queuing

```
router1#config t
Enter configuration commands, one per line.  End with CNTL/Z.
router1(config)#interface serial 0
router1(config-if)#
router1(config-if)# ?
Interface configuration commands:
  access-expression           Build a bridge boolean access expression
  arp                         Set arp type (arpa, probe, snap) or timeout
  backup                      Modify dial-backup parameters
  bandwidth                   Set bandwidth informational parameter
  bridge-group                Transparent bridging interface parameters
  carrier-delay               Specify delay for interface transitions
  cdp                         CDP interface subcommands
  clock                       Configure serial interface clock
  compress                    Set serial interface for compression
  custom-queue-list           Assign a custom queue list to an interface
  dce-terminal-timing-enable  Enable DCE terminal timing
  delay                       Specify interface throughput delay
  description                 Interface specific description
  dialer                      Dial-on-demand routing (DDR) commands
  dialer-group                Assign interface to dialer-list
  down-when-looped            Force looped serial interface down
  dte-invert-txc              Invert transmit clock
  dxi                         ATM-DXI configuration commands
  encapsulation               Set encapsulation type for an interface
  exit                        Exit from interface configuration mode
  fair-queue                  Enable Fair Queuing on an Interface
  frame-relay                 Set frame relay parameters
  full-duplex                 Configure full-duplex operational mode
  half-duplex                 Configure half-duplex and related commands
  help                        Description of the interactive help system
  hold-queue                  Set hold queue depth
  idle-character              Set idle character type
  ignore-dcd                  ignore dcd
  invert-txc                  Invert transmit clock
  ip                          Interface Internet Protocol config commands
  keepalive                   Enable keepalive
  lapb                        X.25 Level 2 parameters (Link Access Procedure,
                              Balanced)
  llc2                        LLC2 Interface Subcommands
  load-interval               Specify interval for load calculation for an
                              interface
  loopback                    Configure internal loopback on an interface
  mac-address                 Manually set interface MAC address
  media-type                  Interface media type
  mtu                         Set the interface Maximum Transmission Unit (MTU)
  netbios                     Use a defined NETBIOS access list or enable
                              name-caching
  no                          Negate a command or set its defaults
  nrzi-encoding               Enable use of NRZI encoding
  ntp                         Configure NTP
  ppp                         Point-to-point protocol
  priority-group              Assign a priority group to an interface
  pulse-time                  Enables pulsing of DTR during resets
  rx-sync-disable             rx sync disable
  serial                      serial interface commands
```

```
  shutdown                  Shutdown the selected interface
  smds                      Modify SMDS parameters
  snapshot                  Configure snapshot support on the interface
  snmp                      Modify SNMP interface parameters
  transmit-interface        Assign a transmit interface to a receive-only
                            interface
  transmitter-delay         Set dead-time after transmitting a datagram
  tx-queue-limit            Configure card level transmit queue limit

router1(config-if)#
router1(config-if)#fair-queue ?
  <1-4096>  Congestive Discard Threshold
  <cr>

router1(config-if)#fair-queue 96 ?
  <cr>

router1(config-if)#fair-queue 96
router1(config-if)#
router1(config-if)#^Z
router1#

router1#
router1#show run
Building configuration...

Current configuration:
!
version 11.0
service udp-small-servers
service tcp-small-servers
!
hostname router1
!
enable secret 5 $1$8gsO$iFE02FDfh3ZbzZ0BujHPO/
enable password 12345
!
!
interface Loopback0
 ip address 172.16.1.1 255.255.255.0
!
interface Loopback1
 ip address 172.16.2.1 255.255.255.0
!
interface Ethernet0
 no ip address
 shutdown
!
interface Ethernet1
 no ip address
 shutdown
!
interface Serial0
 ip address 172.16.3.1 255.255.255.0
 fair-queue 96
 clockrate 56000
!
interface Serial1
 no ip address
 shutdown
!
```

```
router eigrp 20
 network 172.16.0.0
!
!
line con 0
line aux 0
 transport input all
line vty 0 4
 password 12345
 login
!
end
```

```
router1#show ?
  access-expression  List access expression
  access-lists       List access lists
  accounting         Accounting data for active sessions
  aliases            Display alias commands
  arp                ARP table
  async              Information on terminal lines used as router interfaces
  bridge             Bridge Forwarding/Filtering Database [verbose]
  buffers            Buffer pool statistics
  cdp                CDP information
  clock              Display the system clock
  cmns               Connection-Mode networking services (CMNS) information
  compress           Show compression statistics.
  configuration      Contents of Non-Volatile memory
  controllers        Interface controller status
  debugging          State of each debugging option
  dhcp               Dynamic Host Configuration Protocol status
  dialer             Dialer parameters and statistics
  dnsix              Shows Dnsix/DMDP information
  dxi                atm-dxi information
  entry              Queued terminal entries
  flash              System Flash information
  flh-log            Flash Load Helper log buffer
  frame-relay        Frame-Relay information
  history            Display the session command history
  hosts              IP domain-name, lookup style, nameservers, and host table
  interfaces         Interface status and configuration
  ip                 IP information
  line               TTY line information
  llc2               IBM LLC2 circuit information
  logging            Show the contents of logging buffers
  memory             Memory statistics
  ntp                Network time protocol
  ppp                PPP parameters and statistics
  printers           Show LPD printer information
  privilege          Show current privilege level
  processes          Active process statistics
  protocols          Active network routing protocols
  queue              Show queue contents
  queueing           Show queueing configuration
  registry           Function registration information
  reload             Scheduled reload information
  rhosts             Remote-host+user equivalences
  rif                RIF cache entries
  route-map          route-map information
  running-config     Current operating configuration
  sessions           Information about Telnet connections
  smds               SMDS information
  smf                Software MAC filter
  snapshot           Snapshot parameters and statistics
```

```
router1#show queueing ?
  custom    custom queueing list configuration
  fair      fair queueing list configuration
  priority  priority queueing list configuration
  <cr>
```

Copyright : HiTek Solutions Inc.

Configuring Frame Relay & Policy Based Routing

Lab Exercises

Frame Relay configuration example

One cisco 2500 series router will act as a frame relay switch and simulate the "frame relay cloud". The other two routers will be frame relay client routers with an Nt client attached. When working the two client pc's should be able to ping each other and map drives.

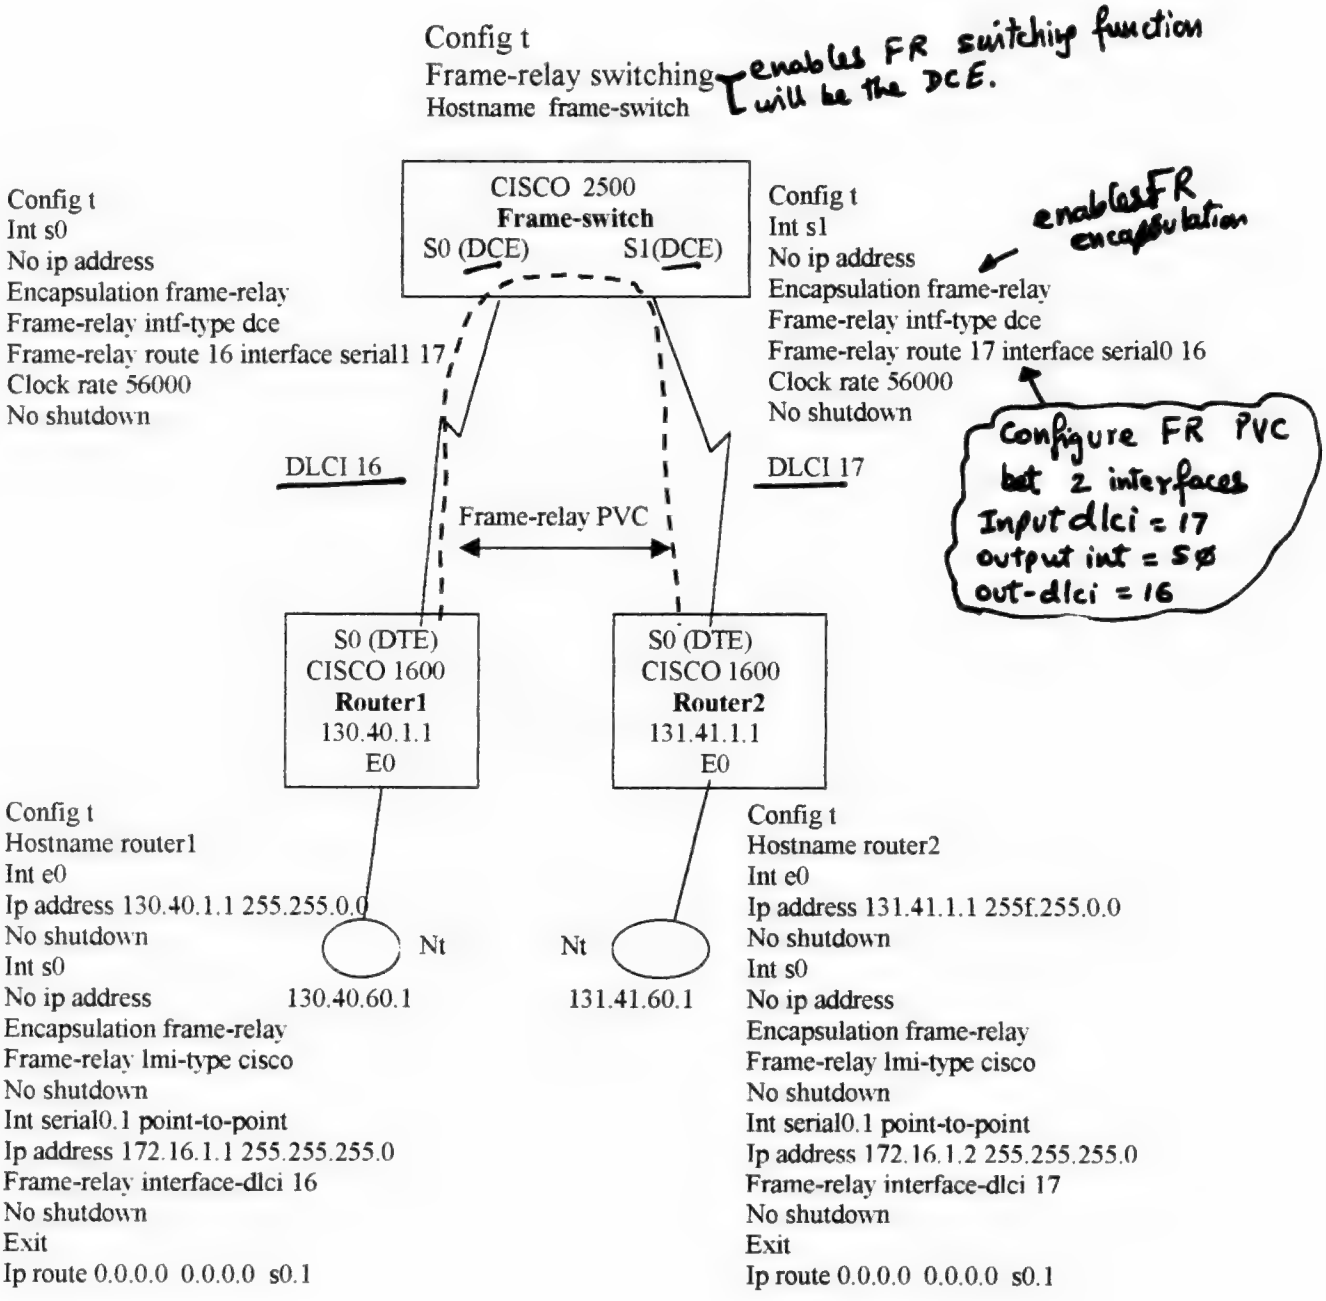

Frame relay switch configuration

```
frame-switch#show run
Building configuration...

Current configuration:
!
version 11.2
!
hostname frame-switch
!
enable secret 5 $1$J/qE$wokIto2s4718DcCFcnpkx0
enable password 12345
!
frame-relay switching
!
interface Ethernet0
 no ip address
 shutdown
!
interface Ethernet1
 no ip address
 shutdown
!
interface Serial0
 no ip address
 encapsulation frame-relay
 clockrate 56000
 frame-relay intf-type dce
 frame-relay route 16 interface Serial1 17
!
interface Serial1
 no ip address
 encapsulation frame-relay
 clockrate 56000
 frame-relay intf-type dce
 frame-relay route 17 interface Serial0 16
!
no ip classless
!
line con 0
line aux 0
line vty 0 4
 password 12345
 login
!
end
```

```
frame-switch#
frame-switch#show int s0
Serial0 is up, line protocol is up
  Hardware is HD64570
  MTU 1500 bytes, BW 1544 Kbit, DLY 20000 usec, rely 255/255, load
1/255
  Encapsulation FRAME-RELAY, loopback not set, keepalive set (10 sec)
  LMI enq sent  2, LMI stat recvd 0, LMI upd recvd 0
  LMI enq recvd 132, LMI stat sent  132, LMI upd sent  0, DCE LMI up
  LMI DLCI 1023  LMI type is CISCO  frame relay DCE
  Broadcast queue 0/64, broadcasts sent/dropped 0/0, interface
broadcasts 0
  Last input 00:00:02, output 00:00:02, output hang never
  Last clearing of "show interface" counters never
  Input queue: 0/75/0 (size/max/drops); Total output drops: 0
  Queueing strategy: weighted fair
  Output queue: 0/64/0 (size/threshold/drops)
     Conversations  0/1 (active/max active)
     Reserved Conversations 0/0 (allocated/max allocated)
  5 minute input rate 0 bits/sec, 0 packets/sec
  5 minute output rate 0 bits/sec, 0 packets/sec
     215 packets input, 14780 bytes, 0 no buffer
     Received 0 broadcasts, 0 runts, 0 giants
     0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored, 0 abort
     185 packets output, 10885 bytes, 0 underruns
     0 output errors, 0 collisions, 50 interface resets
     0 output buffer failures, 0 output buffers swapped out
     39 carrier transitions
     DCD=up  DSR=up  DTR=up  RTS=up  CTS=up
frame-switch#
frame-switch#show int s1
Serial1 is up, line protocol is up
  Hardware is HD64570
  MTU 1500 bytes, BW 1544 Kbit, DLY 20000 usec, rely 255/255, load
1/255
  Encapsulation FRAME-RELAY, loopback not set, keepalive set (10 sec)
  LMI enq sent  1, LMI stat recvd 0, LMI upd recvd 0
  LMI enq recvd 164, LMI stat sent  162, LMI upd sent  0, DCE LMI up
  LMI DLCI 1023  LMI type is CISCO  frame relay DCE
  Broadcast queue 0/64, broadcasts sent/dropped 0/0, interface
broadcasts 0
  Last input 00:00:02, output 00:00:02, output hang never
  Last clearing of "show interface" counters never
  Input queue: 0/75/0 (size/max/drops); Total output drops: 0
  Queueing strategy: weighted fair
  Output queue: 0/64/0 (size/threshold/drops)
     Conversations  0/2 (active/max active)
     Reserved Conversations 0/0 (allocated/max allocated)
  5 minute input rate 0 bits/sec, 0 packets/sec
  5 minute output rate 0 bits/sec, 0 packets/sec
     249 packets input, 16058 bytes, 0 no buffer
     Received 0 broadcasts, 0 runts, 0 giants
     0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored, 0 abort
     248 packets output, 16688 bytes, 0 underruns
     0 output errors, 0 collisions, 17 interface resets
     0 output buffer failures, 0 output buffers swapped out
```

Summary of exchanged data with switch & DTE

```
frame-switch#
frame-switch#show frame-relay ?
  ip       show frame relay IP statistics
  lapf     show frame relay lapf status/statistics
  lmi      show frame relay lmi statistics
  map      Frame-Relay map table
  pvc      show frame relay pvc statistics
  route    show frame relay route
  svc      show frame relay SVC stuff
  traffic  Frame-Relay protocol statistics

frame-switch#show frame-relay lmi
```

or sh frame lmi

Shows status of LMI interface but other directly connected Rtrs

```
LMI Statistics for interface Serial0 (Frame Relay DCE) LMI TYPE = CISCO
  Invalid Unnumbered info 0          Invalid Prot Disc 0
  Invalid dummy Call Ref 0           Invalid Msg Type 0
  Invalid Status Message 0           Invalid Lock Shift 0
  Invalid Information ID 0           Invalid Report IE Len 0
  Invalid Report Request 0           Invalid Keep IE Len 0
  Num Status Enq. Rcvd 135           Num Status msgs Sent 135
  Num Update Status Sent 0           Num St Enq. Timeouts 46

LMI Statistics for interface Serial1 (Frame Relay DCE) LMI TYPE = CISCO
  Invalid Unnumbered info 0          Invalid Prot Disc 0
  Invalid dummy Call Ref 0           Invalid Msg Type 0
  Invalid Status Message 0           Invalid Lock Shift 0
  Invalid Information ID 2           Invalid Report IE Len 0
  Invalid Report Request 0           Invalid Keep IE Len 0
  Num Status Enq. Rcvd 166           Num Status msgs Sent 164
  Num Update Status Sent 0           Num St Enq. Timeouts 3
```

```
frame-switch#
frame-switch#show frame-relay pvc
```

display all PVCs passing through the Rtr

```
PVC Statistics for interface Serial0 (Frame Relay DCE)

DLCI = 16, DLCI USAGE = SWITCHED, PVC STATUS = ACTIVE, INTERFACE =
Serial0
```

shows its a FR switch

```
  input pkts 54             output pkts 54           in bytes 9610
  out bytes 9263            dropped pkts 0           in FECN pkts 0
  in BECN pkts 0            out FECN pkts 0          out BECN pkts 0
  in DE pkts 0              out DE pkts 0
  pvc create time 00:33:35, last time pvc status changed 00:13:57
  Num Pkts Switched 54

PVC Statistics for interface Serial1 (Frame Relay DCE)

DLCI = 17, DLCI USAGE = SWITCHED, PVC STATUS = ACTIVE, INTERFACE =
Serial1

  input pkts 55             output pkts 54           in bytes 9557
  out bytes 9610            dropped pkts 1           in FECN pkts 0
  in BECN pkts 0            out FECN pkts 0          out BECN pkts 0
  in DE pkts 0              out DE pkts 0
  pvc create time 00:32:19, last time pvc status changed 00:22:21
  Num Pkts Switched 54
```

- how many packets were switched.

```
frame-switch#
```

```
frame-switch#
frame-switch#show frame-relay route
```

shows all DLCIs configured on Rtr.

Input Intf	Input Dlci	Output Intf	Output Dlci	Status
Serial0	16	Serial1	17	active
Serial1	17	Serial0	16	active

shows how they are "linked"

```
frame-switch#
```

The following is one of the frame relay client router configurations

```
router2#show run
Building configuration...

Current configuration:
!
version 12.0
service timestamps debug uptime
service timestamps log uptime
no service password-encryption
!
hostname router2
!
enable secret 5 $1$6/10$upCzhjwEEiqAaYQhWFqfd/
enable password 12345
!
ip subnet-zero
!
interface Ethernet0
 ip address 131.41.1.1 255.255.0.0
 no ip directed-broadcast
!
interface Serial0
 no ip address
 no ip directed-broadcast
 encapsulation frame-relay
 frame-relay lmi-type cisco
!
interface Serial0.1 point-to-point
 ip address 172.16.1.2 255.255.255.0
 no ip directed-broadcast
 frame-relay interface-dlci 17
!
ip classless
ip route 0.0.0.0 0.0.0.0 Serial0.1
!
!
line con 0
 transport input none
line 1
line vty 0 4
 password 12345
 login
!
end
```

```
router2#
router2#show ip route
Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B -
BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, * - candidate
default
       U - per-user static route, o - ODR
       T - traffic engineered route

Gateway of last resort is 0.0.0.0 to network 0.0.0.0

     172.16.0.0/24 is subnetted, 1 subnets
C       172.16.1.0 is directly connected, Serial0.1
C    131.41.0.0/16 is directly connected, Ethernet0
S*   0.0.0.0/0 is directly connected, Serial0.1

router2#
```

```
router2#
router2#show int s0
Serial0 is up, line protocol is up
  Hardware is QUICC Serial
  MTU 1500 bytes, BW 1544 Kbit, DLY 20000 usec, rely 255/255, load
1/255
  Encapsulation FRAME-RELAY, loopback not set, keepalive set (10 sec)
  LMI enq sent  112, LMI stat recvd 112, LMI upd recvd 0, DTE LMI up
  LMI enq recvd 0, LMI stat sent  0, LMI upd sent  0
  LMI DLCI 1023  LMI type is CISCO  frame relay DTE
  Broadcast queue 0/64, broadcasts sent/dropped 19/0, interface
broadcasts 0
  Last input 00:00:07, output 00:00:07, output hang never
  Last clearing of "show interface" counters never
  Input queue: 0/75/0 (size/max/drops); Total output drops: 0
  Queueing strategy: weighted fair
  Output queue: 0/1000/64/0 (size/max total/threshold/drops)
     Conversations  0/1/256 (active/max active/max total)
     Reserved Conversations 0/0 (allocated/max allocated)
  5 minute input rate 0 bits/sec, 0 packets/sec
  5 minute output rate 0 bits/sec, 0 packets/sec
     169 packets input, 11158 bytes, 0 no buffer
     Received 112 broadcasts, 0 runts, 0 giants, 0 throttles
     8 input errors, 0 CRC, 8 frame, 0 overrun, 0 ignored, 0 abort
     196 packets output, 14792 bytes, 0 underruns
     0 output errors, 0 collisions, 18 interface resets
     0 output buffer failures, 0 output buffers swapped out
     0 carrier transitions
     DCD=up  DSR=up  DTR=up  RTS=up  CTS=up

router2#
router2#
router2#show frame-relay ?
  ip              show frame relay IP statistics
  lapf            show frame relay lapf status/statistics
  lmi             show frame relay lmi statistics
  map             Frame-Relay map table
  pvc             show frame relay pvc statistics
  qos-autosense   show frame relay qos-autosense information
  route           show frame relay route
  svc             show frame relay SVC stuff
  traffic         Frame-Relay protocol statistics


router2#
router2#show frame-relay lmi

LMI Statistics for interface Serial0 (Frame Relay DTE) LMI TYPE = CISCO
  Invalid Unnumbered info 0          Invalid Prot Disc 0
  Invalid dummy Call Ref 0           Invalid Msg Type 0
  Invalid Status Message 0           Invalid Lock Shift 0
  Invalid Information ID 0           Invalid Report IE Len 0
  Invalid Report Request 0           Invalid Keep IE Len 0
  Num Status Enq. Sent 115           Num Status msgs Rcvd 115
  Num Update Status Rcvd 0           Num Status Timeouts 0
```

```
router2#
router2#show frame-relay map
Serial0 (up): ip 0.0.0.0 dlci 17(0x11,0x410)
              broadcast,
              CISCO, status defined, active
Serial0.1 (up): point-to-point dlci, dlci 17(0x11,0x410), broadcast
          status defined, active

router2#

router2#show frame-relay pvc

PVC Statistics for interface Serial0 (Frame Relay DTE)

DLCI = 17, DLCI USAGE = LOCAL, PVC STATUS = ACTIVE, INTERFACE =
Serial0.1

  input pkts 56             output pkts 54           in bytes 9248
  out bytes 9260            dropped pkts 0           in FECN pkts 0
  in BECN pkts 0            out FECN pkts 0          out BECN pkts 0
  in DE pkts 0              out DE pkts 0
  out bcast pkts 20          out bcast bytes 5880
  pvc create time 00:19:21, last time pvc status changed 00:19:21
router2#
```

Notice that router2 can ping router1 thru the frame relay network !!!

```
router2#
router2#ping 130.40.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 130.40.1.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 72/73/76 ms

router2#
router2#traceroute 130.40.1.1

Type escape sequence to abort.
Tracing the route to 130.40.1.1

  1 172.16.1.1 40 msec 60 msec *
router2#
router2#
```

Policy-based routing

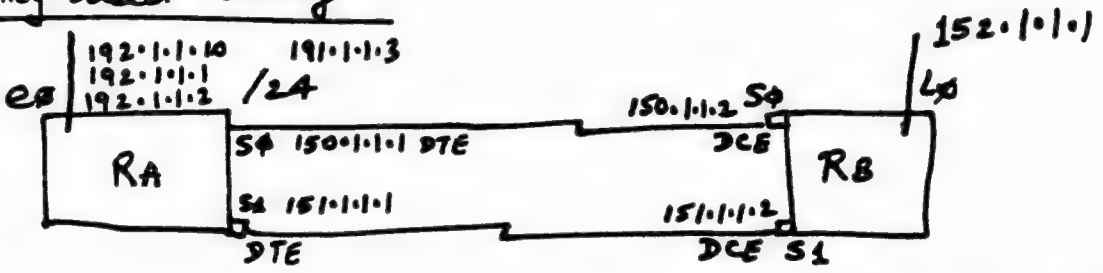

GOAL → TRAFFIC FROM 192.1.1.1 will be routed OUT on S0
on RTRA → TRAFFIC FROM 192.1.1.2 will be routed OUT on S1

192.1.1.1 & 192.1.1.3 are test pts
192.1.1.2

RA has multiple secondary interfaces on e0. We will create/apply a route-map ip policy called lab1 to e0 on RA which will set the next hop to S0 interface for packets from 192.1.1.1 & " " " " " S1 " " " " 192.1.1.2.

RA

```
hostname RA
int e0
 ip address 192.1.1.1 255.255.255.0 secondary
 ip address 192.1.1.2 255.255.255.0 secondary
 ip address 192.1.1.3 255.255.255.0 secondary
 ip address 192.1.1.10 255.255.255.0
 ip policy route-map lab1
 no keepalive

int s0
 ip address 150.1.1.1 255.255.255.0
 no fair-queue

int s1
 ip address 151.1.1.1 255.255.255.0

router rip
 network 150.1.0.0
 network 151.1.0.0
 network 192.1.1.0

ip local policy route-map lab1   <- enables a route policy
 no ip classless
 access-list 1 permit 192.1.1.1
 access-list 2 permit 192.1.1.2
 route-map lab1 permit 10   <- 10 = item # in route map called lab1
  match ip address 1   <- match criteria from access list 1
  set interface S0
 route-map lab1 permit 20   <- 20 is item # in route map called lab1
  match ip address 2
  set inteface S1
```

RB

```
int loopback 0
 ip address 152.1.1.1 255.255.255.0
int e0
 no ip address
 shutdown
router rip
 network 152.1.1.0
 network 150.1.1.0
 network 151.1.1.0
```

```
int s0
 Clock rate 56000
 ip address 150.1.1.2 255.255.255.0
 no shut
int s1
 clockrate 56000
 ip address 151.1.1.2 255.255.255.0
 no shut
```

Now @ RA

```
# traceroute ip
or traceroute 152.1.1.1
& traceroute 150.1.1.2
& traceroute 151.1.1.2
```

```
debug ip rip
ping <enter>
 Specify target = 152.1.1.1
 Specify source = 192.1.1.2
 Yes for extended cmds.
 No for rest
```

---

on RA

In the dia. we can set up something like this.

```
ip local policy route-map lab1
 no ip classless
 access-list 101 permit tcp any any eq www
 access-list 102 permit tcp any any eq telnet
 route-map lab1 permit 10
 match ip address 101
 set ip next-hop 150.1.1.2
 route-map lab1 permit 20
 match ip address 102
 set ip next-hop 151.1.1.2
```

(Policy-based routing)

① What do you accomplish ?

→ Send http traffic via S0

→ send telnet traffic via S1.

DDR

```
config t
int s0
backup interface bri0
backup delay 5 60
```

ISDN will be backup for S0

- wait for 5 sec. after S0 fails
- After ISDN connection is being used & S0 is back on line, wait for 60 sec before shutting ISDN down

```
config t
int s0
backup int bri0
backup load 80 30
```

For S0 if usage is 80% start ISDN.

For S0 when bandwidth usage is dn to 30% shut dn ISDN.

Copyright : HiTek Solutions, Inc

Basic OSPF router configuration in a single area

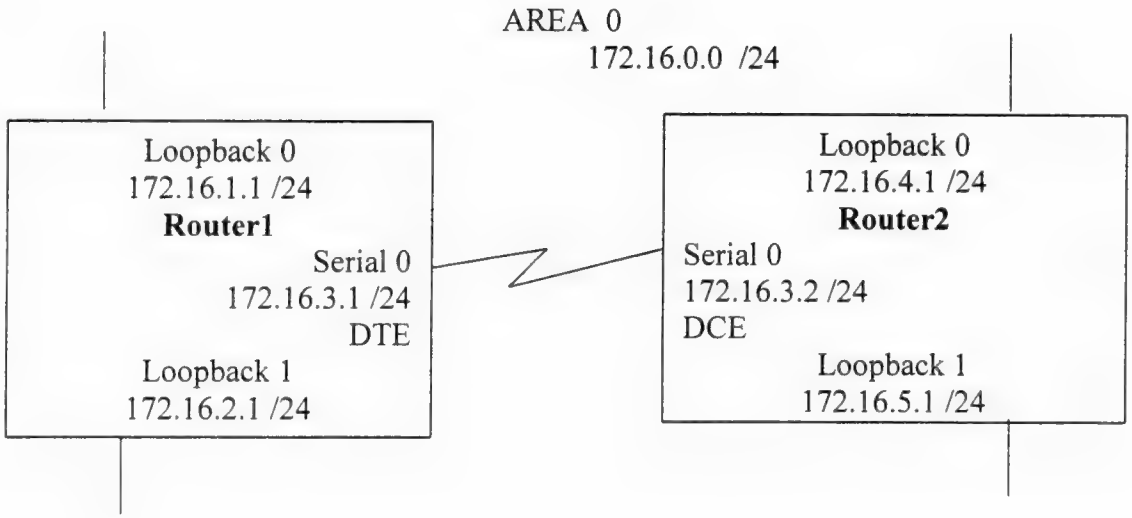


```
Config t
Hostname router1
Ip routing
Int loopback0
Ip address 172.16.1.1 255.255.255.0
Int loopback1
Ip address 172.16.2.1 255.255.255.0
Int serial0
Ip address 172.16.3.1 255.255.255.0
No shutdown

Router ospf 10
Network 172.16.1.1 0.0.0.0 area 0
Network 172.16.2.1 0.0.0.0 area 0
Network 172.16.3.1 0.0.0.0 area 0
```

```
Config t
hostname router2
ip routing
int loopback0
ip address 172.16.4.1 255.255.255.0
int loopback1
ip address 172.16.5.1 255.255.255.0
int serial0
ip address 172.16.3.2 255.255.255.0
no shutdown
Clock rate 56000

router ospf 10
network 172.16.4.1 0.0.0.0 area 0
network 172.16.5.1 0.0.0.0 area 0
network 172.16.3.2 0.0.0.0 area 0
```

OR ANOTHER WAY

```
Router ospf 10
Network 172.16.0.0  0.0.255.255 area 0
```

```
router ospf 10
network 172.16.0.0 0.0.255.255 area 0
```

Basic OSPF router configuration for multiple area's

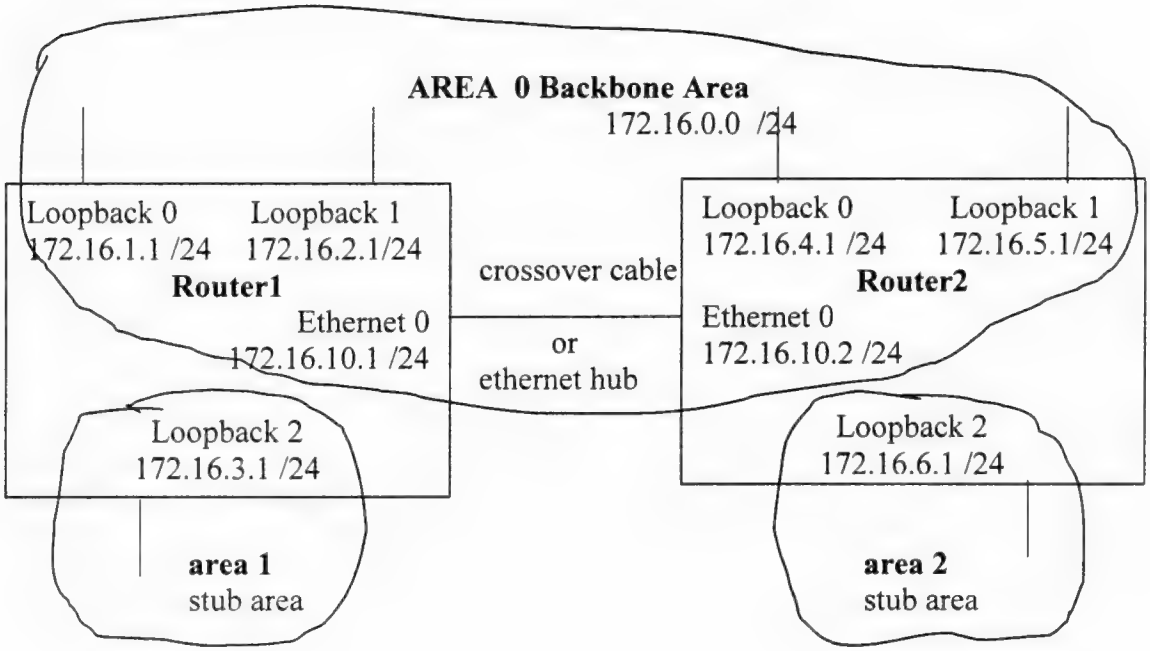


```
Config t
Hostname router1
Ip routing
Int loopback0
Ip address 172.16.1.1 255.255.255.0
Int loopback1
Ip address 172.16.2.1 255.255.255.0
Int loopback 2
Ip address 172.16.3.1 255.255.255.0
Int ethernert 0
172.16.10.1 255.255.255.0
no shutdown

Router ospf 20
Network 172.16.3.1 0.0.0.0 area 1
Network 172.16.0.0 0.0.255.255 area 0
Area 1 stub
```

```
Config t
hostname router2
ip routing
int loopback0
ip address 172.16.4.1 255.255.255.0
int loopback1
ip address 172.16.5.1 255.255.255.0
int loopback 2
ip address 172.16.6.1 255.255.255.0
interface ethernet 0                    ip address
ip address 172.16.10.2 255.255.255.0 no shutdown

router ospf 20
network 172.16.6.1 0.0.0.0 area 2
network 172.16.0.0 0.0.255.255 area 0
area 2 stub
```

Basic EIGRP router configuration

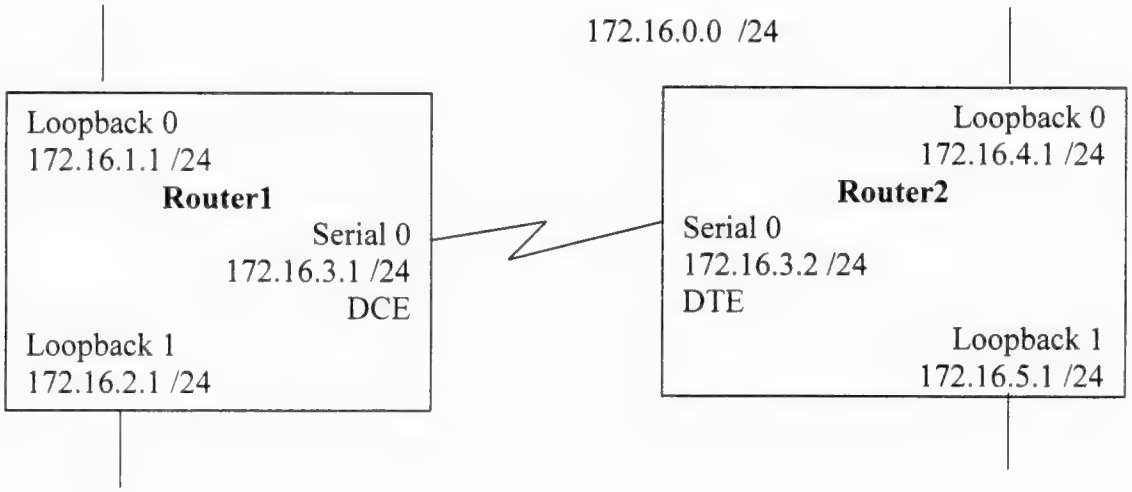

Config t
Hostname router1
Ip routing
Int loopback0
Ip address 172.16.1.1 255.255.255.0
Int loopback1
Ip address 172.16.2.1 255.255.255.0
Int serial0
Ip address 172.16.3.1 255.255.255.0
No shutdown
Clock rate 56000

Router eigrp 20
Network 172.16.1.0
Network 172.16.2.0
Network 172.16.3.0

Config t
hostname router2
ip routing
int loopback0
ip address 172.16.4.1 255.255.255.0
int loopback1
ip address 172.16.5.1 255.255.255.0
int serial0
ip address 172.16.3.2 255.255.255.0
no shutdown

router eigrp 20
network 172.16.4.0
network 172.16.5.0
network 172.16.3.0

Basic IPX router configuration

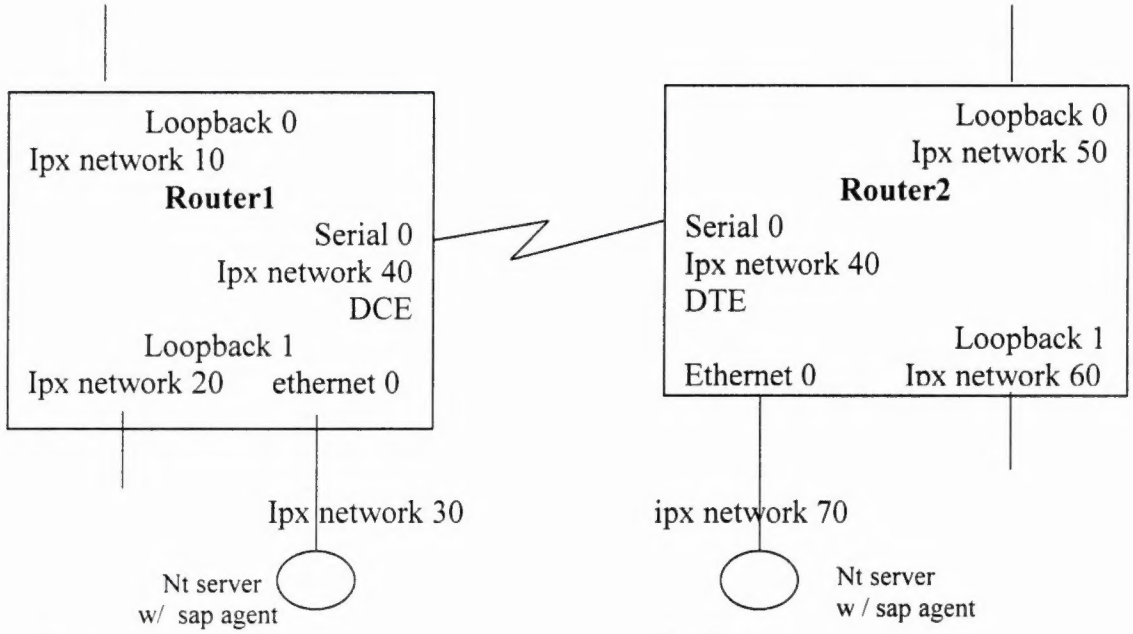

Config t
Hostname router1
Ipx routing
Int loopback0
Ipx network 10
Int loopback1
Ipx network 20
Int serial0
Ipx network 40
No shutdown
Clock rate 56000
Int ethernet 0
Ipx network 30 encapsulation sap
No shutdown

Ipx router rip
Network 10
Network 20
Network 30
Network 40

Config t
hostname router2
ipx routing
int loopback0
ipx network 50
int loopback1
ipx network 60
int serial0
ipx network 40
no shutdown

int ethernet 0
ipx network 70 encapsulation sap
no shutdown

ipx router rip
network 40
network 50
network 60
network 70